arXiv:1612.01619v1 [stat.OT] 6 Dec 2016

High-dimensional nonparametric monotone function estimation using BART

Hugh A. Chipman, Edward I. George, Robert E. McCulloch, Thomas S. Shively *

December 7, 2016

*Hugh Chipman is Professor, Department of Mathematics and Statistics, Acadia University, Wolfville, Nova Scotia, B4P 2R6, hugh.chipman@acadiau.ca. Edward I. George is Professor of Statistics, Department of Statistics, The Wharton School, University of Pennsylvania, 3730 Walnut St, 400 JMHH, Philadelphia, PA 19104-6304, edgeorge@wharton.upenn.edu. Robert E. McCulloch is Professor of Statistics, The School of Mathematical and Statistical Sciences, Arizona State University, Tempe, AZ, 85281, robert.mcculloch@asu.edu. Thomas S. Shively is Professor of Statistics, Department of Information, Risk, and Operations Management, University of Texas at Austin, 1 University Station, B6500, Austin, TX 78712-1175, Tom.Shively@mccombs.utexas.edu. This work was supported by the grant DMS-1406563 from the National Science Foundation.

Abstract

For the estimation of a regression relationship between Y and a large set of potential predictors $x_1, \ldots, x_p$, the flexible nature of a nonparametric approach such as BART (Bayesian Additive Regression Trees) allows for a much richer set of possibilities than a more restrictive parametric approach. However, it may often occur that subject matter considerations suggest the relationship will be monotone in one or more of the predictors. For such situations, we propose monotone BART, a constrained version of BART that uses the monotonicity information to improve function estimation without the need of using a parametric form. Imposing monotonicity, when appropriate, results in (i) function estimates that are smoother and more interpretable, (ii) better out-of-sample predictive performance, (iii) less uncertainty, and (iv) less sensitivity to prior choice. While some of the key aspects of the unconstrained BART model carry over directly to monotone BART, the imposition of the monotonicity constraints necessitates a fundamental rethinking of how the model is implemented. In particular, in the original BART algorithm, the Markov Chain Monte Carlo algorithm relied on a conditional conjugacy that is no longer available in a high-dimensional, constrained space.

KEY WORDS: Bayesian nonparametrics, high-dimensional regression, ensemble model, MCMC algorithm, monotonicity constraints

Contents

1 Introduction **1**

2 A Monotone Sum-of-Trees Model **2**

3 A Constrained Regularization Prior **7**

3.1 The T_j Prior 8

3.2 The σ Prior 9

3.3 The $M_j \mid T_j$ Prior 9

3.3.1 The Choice of m 11

4 Simulation from the Constrained Posterior **11**

4.1 Backfitting MCMC for Constrained Regression Trees 11

4.2 A New Localized MH Algorithm 12

4.3 Implementation of the Localized MH Algorithm 16

5 Examples **19**

5.1 One-Dimensional Simulated Example 20

5.2 Five-Dimensional Simulated Example 21

5.3 The House Price Example 23

5.4 The Car Price Example 29

5.5 The Stock Returns Example 35

6 Conclusion **42**

1 Introduction

Suppose one would like to learn how Y depends on a vector of potential predictors $x = (x_1, \ldots, x_p)$ when very little prior information is available about the form of the relationship. With only very weak assumptions, the Bayesian nonparametric approach BART (Bayesian Additive Regression Trees) can quickly discover the nature of this relationship; see Chipman, George, and McCulloch (2010), hereafter CGM10. More precisely, based only on the assumption that

$$Y = f(x) + \epsilon, \qquad \epsilon \sim N(0, \sigma^2), \tag{1.1}$$

BART can quickly obtain full posterior inference for the unknown regression function,

$$f(x) = E(Y \mid x) \tag{1.2}$$

and the unknown variance σ^2. BART also provides predictive inference as well as model-free variable selection and interaction detection, see Chipman, George, and McCulloch (2013), Bleich et al. (2014), and Kapelner and Bleich (2016).

The main goal of this paper is the development of monotone BART (hereafter mBART), a constrained version of BART that restricts attention to regression functions $f(x)$ that are monotone in any predesignated subset of the components of x. Such monotonicity constraints often arise naturally from subject matter considerations. For example, in one of our illustrative data sets, the Y variable of interest is the price of a house. One of the components of the explanatory variable x is the size of the house. It seems reasonable to restrict our search for the function f to functions such that bigger houses sell for more, all other things equal.

There is a rich literature on monotone function estimation in both the frequentist and Bayesian paradigms. Frequentist methods to estimate univariate monotone functions include Barlow et al. (1972), Mammen (1991), Ramsay (1998) and Kong and Eubank (2006) while Bayesian methods include Lavine and Mockus (1995), Holmes and Heard (2003), Neelon and Dunson (2004), Shively, Sage, Walker (2009), and Shively, Walker, Damien (2011). However, these methods are difficult to implement for high-dimensional multivariate function estimation because they are built on basis elements that are fundamentally low-dimensional objects. Saarela and Arjas (2011) and Lin and Dunson (2014) develop methods to estimate monotone multivariate functions but they become computationally intensive as the number of predictor variables increase. As we will show, mBART imposes no restrictions on f beyond the monotonicity constraints while at the same time, it can easily handle high-dimensional

data. The reason is that mBART is built on a sum-of-trees approximation to f and is therefore composed of multivariate basis elements.

The extension of BART to our monotonically constrained setting requires two basic innovations. First, it is necessary to develop general constraints for regression trees to be monotone in any predesignated set of coordinates. Under these constraints, the monotonicity of the complete sum-of-trees approximation follows directly. The second innovation requires a new approach for MCMC posterior computation. Whereas the original BART formulation allowed straightforward marginalization over regression tree parameters using a conditional conjugacy argument, the constrained trees formulation requires a more nuanced approach because the conjugacy argument no longer applies.

The outline of the paper is as follows. In Section 2, we describe in detail the constrained sum-of-trees model used for monotone function estimation. Section 3 discusses the regularization prior for the trees while section 4 describes the new MCMC algorithm required to implement mBART. Section 5 provides simulation results for one-dimensional and five-dimensional function estimation as well as three examples using house price, car price and stock returns data. Section 6 contains some concluding remarks.

2 A Monotone Sum-of-Trees Model

The essence of BART is a sum-of-trees model approximation of the relationship between y and x in (1.1);

$$Y = \sum_{j=1}^{m} g(x; T_j, M_j) \, + \, \epsilon, \qquad \epsilon \sim N(0, \sigma^2), \tag{2.1}$$

where each T_j is a binary regression tree with a set M_j of associated terminal node constants μ_{ij}, and $g(x; T_j, M_j)$ is the function which assigns $\mu_{ij} \in M_j$ to x according to the sequence of decision rules in T_j. These decision rules are binary partitions of the predictor space of the form $\{x \leq a\}$ vs $\{x > a\}$ where the splitting value a is in the range of x. (A clarifying example of how g works appears in Figure 1 below and is described later in this section). When $m = 1$, (2.1) reduces to the single tree model used by Chipman et al. (1998) for Bayesian CART.

Under (2.1), $E(Y \,|\, x)$ is the sum, over trees $T_1, \ldots, T_m$, of all the terminal node μ_{ij}'s assigned to x by the $g(x; T_j, M_j)$'s. As the μ_{ij} can take any values it is easy to see that the sum-of-trees model (2.1) is a very flexible representation capable of representing a wide class of

functions from R^n to R, especially when the number of trees m is large. Composed of many simple functions from R^p to R, namely the $g(x; T_j, M_j)$, the sum-of-trees representation is much more manageable than a representation with more complicated basis elements such as multidimensional wavelets or multidimensional splines. And because each tree function g is invariant to monotone transformations of x (aside from the splitting value), standardization choices are not needed for the predictors.

Key to the construction of monotone BART are the conditions under which the underlying sum-of-trees function $\sum_{j=1}^{m} g(x; T_j, M_j)$ will satisfy the following precise definition of a multivariate monotone function.

Definition: For a subset S of the coordinates of $x \in R^n$, a function $f : R^n \to R$ is said to be *monotone in* S if for each $x_i \in S$ and all values of x, f satisfies

$$f(x_1, \ldots, x_i + \delta, \ldots, x_p) \geq f(x_1, \ldots, x_i, \ldots, x_p), \tag{2.2}$$

for all $\delta > 0$ (f is nondecreasing), or for all $\delta < 0$ (f is nonincreasing).

Clearly, a sum-of-trees function will be monotone in S whenever each of the component trees is monotone in S. Thus it suffices to focus on the conditions for a single tree function $g(x; T, M)$ to be monotone in S. As we'll see, this will only entail providing constraints on the set of terminal node constants M; constraints determined by the tree T.

We illustrate these concepts with the bivariate monotone tree function in Figure 1. This tree has six terminal nodes, labeled 4,10,11,12,13, and 7. The labels follow the standard tree node labeling scheme where the top node is labeled 1 and any non-terminal node with label j has a left child with label $2j$ and a right child with label $2j + 1$. Beginning at the top node, each $x = (x_1, x_2)$ is assigned to subsequent nodes according to the sequence of splitting rules it meets. This continues until x reaches a terminal node where $g(x; T, M)$ assigns the designated value of μ from the set M. For example, with this choice of (T, M), $g(x; T, M) = 3$ when $x = (.6, .4)$.

Alternative views of the function in Figure 1 are depicted in Figure 2. Figure 2a shows the partitions of the x space induced by T. The terminal node regions, R_4,R_{10},R_{11},R_{12},R_{13},R_7, correspond to the six similarly labeled terminal nodes of T. Figure 2b shows $g(x; T, M)$ as a simple step function which assigns a level $\mu \in M$ to each terminal node region. From Figure 2b, it is clear that for any $x = (x_1, x_2)$, moving x to $(x_1 + \delta, x_2)$ or to $(x_1, x_2 + \delta)$ can only increase g for $\delta > 0$. Thus, in the sense of our definition, this $g(x; T, M)$ is monotone in both x_1 and x_2.

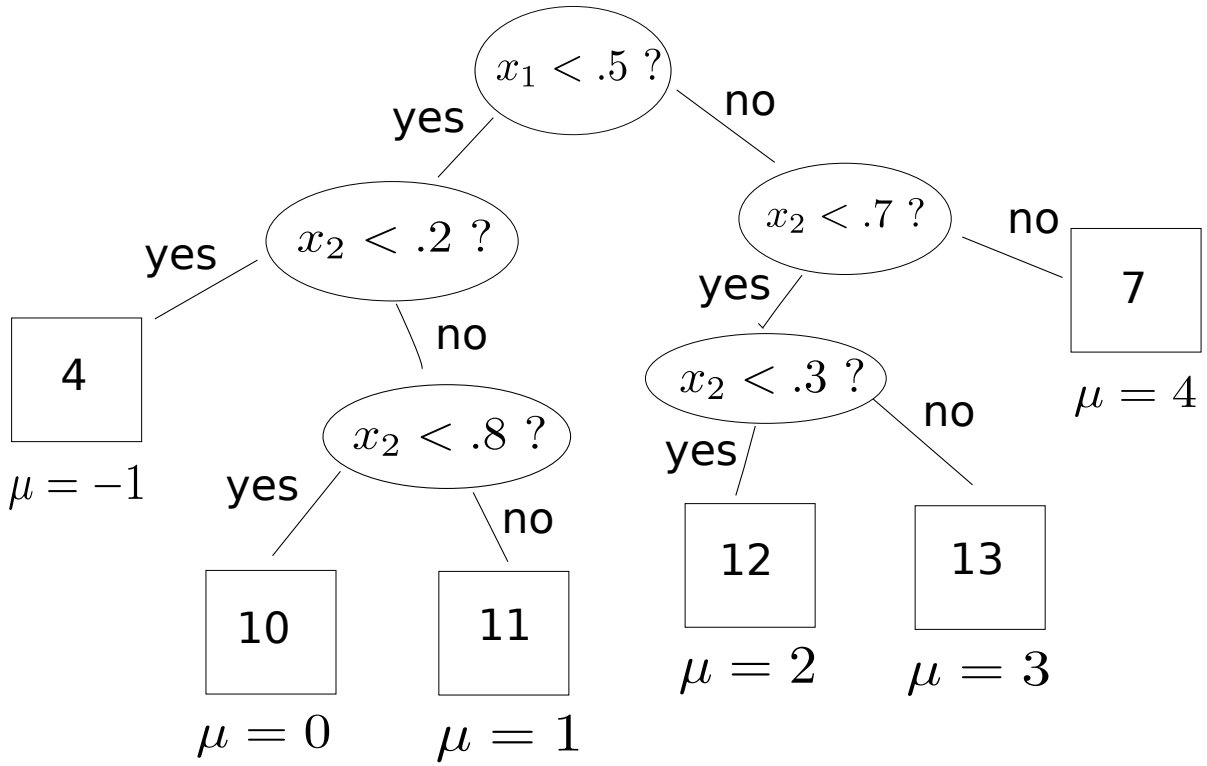

Figure 1: A bivariate, monotone regression tree T with 6 terminal nodes. Intermediate nodes are labeled with their splitting rules. Terminal nodes (bottom leaf nodes) are labeled with their node number. Below each terminal node is the value of $\mu \in M$ assigned to x by $g(x; T, M)$.

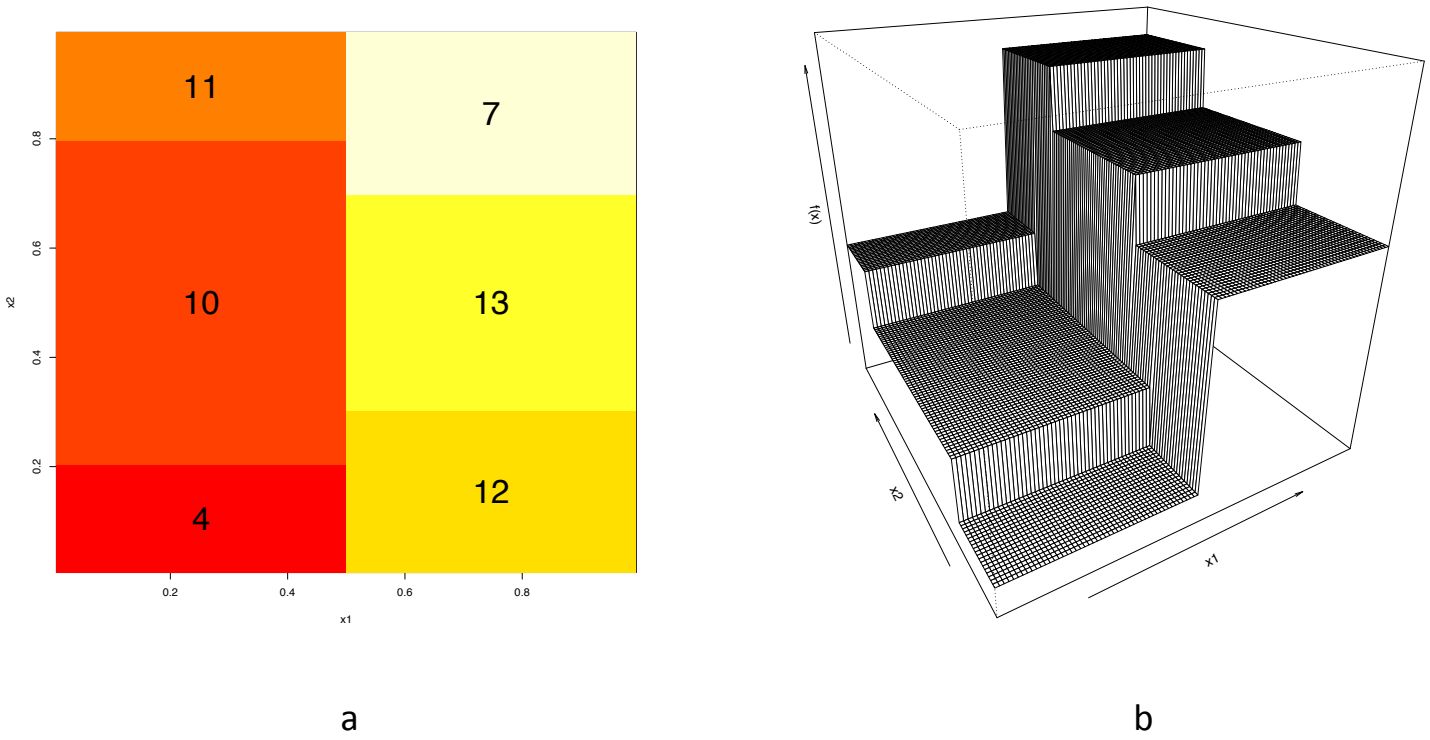

Figure 2: Two alternative views of the bivariate single tree model in Figure 1. (a) The six regions R_4,R_{10},R_{11},R_{12},R_{13},R_7, corresponding to the terminal nodes 4,10,11,12,13,7. (b) The levels of the regions assigned by the step function $g(x; T, M)$.

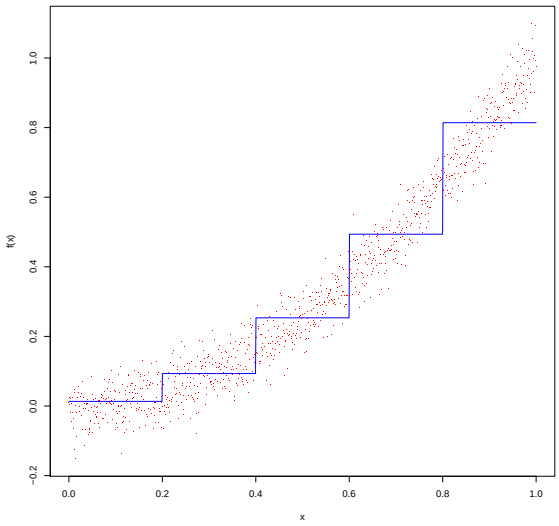

Figure 3: A monotone, univariate tree function $g(x; T, M)$.

To see the essence of what is needed to guarantee the monotonicity of a tree function, consider the very simple case of a monotone $g(x; T, M)$ when T is a function of $x = x_1$ only, as depicted in Figure 3. Each level region of g corresponds to a terminal node region in x_1 space, which is simply an interval whenever g is a univariate function. For each such region, consider the adjoining region with larger values of x_1, which we refer to as an above-neighbor region, and the adjoining region with smaller values of x_1, which we refer to as a below-neighbor region. End regions will only have single neighboring regions. To guarantee (nondecreasing) monotonicity, it suffices to constrain the μ level assigned to each terminal node region to be less than or equal to the μ level of its above-neighbor region, and to be greater than or equal to the μ level of its below-neighbor region.

To apply these notions to a bivariate tree function $g(x; T, M)$ as depicted in Figures 1 and 2, we will simply say that rectangular regions are neighboring if they have boundaries which are adjoining in any of the coordinates. Furthermore, a region R_k will be called an above-neighbor of a region R_{k^*} if the lower adjoining boundary of R_k is the upper adjoining boundary of R_{k^*}. A below-neighbor is defined similarly. For example, in Figure 2a, R_7 is an above-neighbor of R_{10}, R_{11} and R_{13}; and R_{10} and R_{12} are below-neighbors of R_{13}.

Note that R_4 and R_{13} are not neighbors. We will say the R_4 and R_{13} are *separated* because the x_2 upper boundary of R_4 is less than the x_2 lower boundary of R_{13}. For a small enough step size δ, it is impossible to get from R_4 to R_{13} by changing any x_i by δ so that the mean level of one does not constrain the mean level of the other.

To make these definitions precise for a d-dimensional tree T (a function of $x = (x_1, \ldots, x_d)$),

we note that each terminal node region of T will be a rectangular region of the form

$$R_k = \{x : x_i \in [L_{ik}, U_{ik}), i = 1, \ldots, d\}, \tag{2.3}$$

where the interval $[L_{ik}, U_{ik})$ for each x_i is determined by the sequence of splitting rules leading to R_k.

We say that R_k is separated from R_{k^*} if $U_{ik} < L_{ik^*}$ or $L_{ik} > U_{ik^*}$ for some i. In Figure 2a, R_{13} is separated from R_4 and R_{11}.

If R_k and R_{k^*} are not separated, R_k will be said to be an above-neighbor of R_{k^*} if $L_{ik} = U_{ik^*}$ for some i, and it will be said to be a below-neighbor of R_{k^*} if $U_{ik} = L_{ik^*}$ for some i. Note that any terminal node region may have several above-neighbor and below-neighbor regions. R_{13} has below neighbors R_{10} and R_{12} and above neighbor R_7.

The constraints on the μ levels under which $g(x; T, M)$ will be monotone are now straightforward to state.

Constraint Conditions for Tree Monotonicity: A tree function $g(x; T, M)$ will be monotone if the μ level of each of its terminal node regions is
(a) less than or equal to the minimum level of all of its above-neighbor regions, and
(b) greater than or equal to the maximum level of all of its below-neighbor regions.

The function g will be monotone in S if the neighboring regions satisfy (a) and (b) for all the coordinates in S (rather than all coordinates).

As we'll see in subsequent sections, an attractive feature of these conditions is that they dovetail perfectly with the nature of our iterative MCMC simulation calculations. At each step there, we simulate one terminal node level at time conditionally on all the other node levels, so imposing the constraints is straightforward. This avoids the need to simultaneously constrain all the levels at once.

3 A Constrained Regularization Prior

The mBART model specification is completed by putting a constrained regularization prior on the parameters, $(T_1, M_1), \ldots, (T_m, M_m)$ and σ, of the sum-of-trees model (2.1). Essentially a modification of the original BART prior formulation to accommodate designated monotone

constraints, we follow CGM10 and proceed by restricting attention to priors of the form

$$\begin{aligned} p((T_1, M_1), \ldots, (T_m, M_m), \sigma) &= \left[\prod_j p(T_j, M_j)\right] p(\sigma) \\ &= \left[\prod_j p(M_j \mid T_j)\, p(T_j)\right] p(\sigma). \end{aligned} \tag{3.1}$$

Under such priors, the tree components $(T_1, M_1), \ldots, (T_m, M_m)$ are independent of each other and of σ.

As discussed in the previous section, a sum-of-trees function $\sum_{j=1}^m g(x; T_j, M_j)$ is guaranteed to be monotone whenever each of the trees $g(x; T_j, M_j)$ is monotone in the sense of (2.2). Thus, it suffices to restrict the support of $p(M_j \mid T_j)$ to μ_{ij} values which satisfy the Monotonicity Constraints (a) and (b) from Section 2. For this purpose, let C be the set of all (T, M) which satisfy these monotonicity constraints, namely

$$C = \{(T, M) : g(x; T, M) \text{ is monotone in } x_i \in S\}. \tag{3.2}$$

These constraints are then incorporated into the prior by constraining the CGM10 BART independence form $p(M_j \mid T_j) = \prod_i p(\mu_{ij} \mid T_j)$ to have support only over C,

$$p(M_j \mid T_j) \propto \left[\prod_{i=1}^{b_j} p(\mu_{ij} \mid T_j)\right] \chi_C(T_j, M_j). \tag{3.3}$$

Here b_j is the number of bottom (terminal) nodes of T_j, and $\chi_C(\cdot) = 1$ on C and $= 0$ otherwise.

In the next three sections we discuss the choice of priors $p(T_j)$, $p(\sigma)$, and $p(\mu_{ij} \mid T_j)$. These priors will have the same form as in CGM10, but in some cases the monotonicity constraint will motivate modifications to our choices for the hyper parameters.

3.1 The T_j Prior

The tree prior $p(T_j)$ is specified by three aspects: (i) the probability of a node having children at depth d $(= 0, 1, 2, \ldots)$ is

$$\alpha(1 + d)^{-\beta}, \qquad \alpha \in (0, 1), \beta \in [0, \infty), \tag{3.4}$$

(ii) the uniform distribution over available predictors for splitting rule assignment at each interior node, and (iii) the uniform distribution on the discrete set of available splitting

values for the assigned predictor at each interior node. This last choice has the appeal of invariance under monotone transformations of the predictors.

Because we want the regularization prior to keep the individual tree components small, especially when m is set to be large, we typically recommend the defaults $\alpha = .95$ and $\beta = 2$ in (3.4) in the unconstrained case. With this choice, simulation of tree skeletons directly from (i) shows us that trees with 1, 2, 3, 4, and ≥ 5 terminal nodes will receive prior probabilities of about 0.05, 0.55, 0.28, 0.09, and 0.03, respectively.

Discussion of the choice of α and β in the constrained case is deferred to the end of Section 4.3 since our choices are motivated by details of the Markov Chain Monte Carlo algorithm for posterior computation.

3.2 The σ Prior

For $p(\sigma)$, we use the (conditionally) conjugate inverse chi-square distribution $\sigma^2 \sim \nu\,\lambda/\chi^2_\nu$. To guide the specification of the hyperparameters ν and λ, we recommend a data-informed approach in order to assign substantial probability to the entire region of plausible σ values while avoiding overconcentration and overdispersion. This entails calibrating the prior degrees of freedom ν and scale λ using a "rough data-based overestimate" $\hat{\sigma}$ of σ.

The two natural choices for $\hat{\sigma}$ are (1) the "naive" specification, in which we take $\hat{\sigma}$ to be the sample standard deviation of Y (or some fraction of it), or (2) the "linear model" specification, in which we take $\hat{\sigma}$ as the residual standard deviation from a least squares linear regression of Y on the original x's. We then pick a value of ν between 3 and 10 to get an appropriate shape, and a value of λ so that the qth quantile of the prior on σ is located at $\hat{\sigma}$, that is $P(\sigma < \hat{\sigma}) = q$. We consider values of q such as 0.75, 0.90 or 0.99 to center the distribution below $\hat{\sigma}$. For automatic use, we recommend the default setting $(\nu, q) = (3, 0.90)$ which tends to avoid extremes. Alternatively, the values of (ν, q) may be chosen by cross-validation from a range of reasonable choices. This choice is exactly as in CGM10.

3.3 The $M_j \,|\, T_j$ Prior

For the choice of $p(\mu_{ij} \,|\, T_j)$ in (3.3), we adopt a normal form as in BART. However, here we use different choices of the prior variance depending on whether C in (3.2) imposes

constraints on μ_{ij}. For μ_{ij} unconstrained by C, we set

$$p(\mu_{ij} \,|\, T_j) = \phi_{\mu_\mu, \sigma_\mu}, \tag{3.5}$$

the same $N(\mu_\mu, \sigma_\mu^2)$ density used by CGM10 for BART, whereas for μ_{ij} constrained by C, we set

$$p(\mu_{ij} \,|\, T_j) = \phi_{\mu_\mu, c\sigma_\mu}, \tag{3.6}$$

the $N(\mu_\mu, c\sigma_\mu^2)$ density with $c = \frac{\pi}{\pi - 1} \approx 1.4669$.

To motivate our prior choice in (3.6), consider a simple tree with just two terminal node means μ_1 and μ_2 constrained to satisfy $\mu_1 \leq \mu_2$. Under (3.6), the joint distribution of μ_1 and μ_2 is

$$p(\mu_1, \mu_2) \propto \phi_{\mu_\mu, c\sigma_\mu}(\mu_1) \phi_{\mu_\mu, c\sigma_\mu}(\mu_2)\, \chi_{\{\mu_1 \leq \mu_2\}}(\mu_1, \mu_2). \tag{3.7}$$

In this case, the marginal distributions of μ_1 and μ_2 are skew normal distributions (Azzalini 1985) with variances equal to σ_μ^2, and means $\mu_\mu - c\sigma_\mu/\sqrt{\pi}$ and $\mu_\mu + c\sigma_\mu/\sqrt{\pi}$, respectively. Thus, the marginal prior variances of μ_1 and μ_2 are identical to the marginal variances of the unconstrained means. This equality helps to balance the prior effects across predictors and facilitates the calibrated specification of σ_μ described below. Of course, it will be the case that some means μ_{ij} may be further constrained when they occur deeper down the tree, thereby further reducing their prior variance. Although additional small prior adjustments can be considered for such cases, we view them as relatively unimportant because the vast majority of BART trees will be small with at most one or two constraints. Thus, we recommend the prior (3.6) for all μ_{ij} which are constrained.

To guide the specification of the hyperparameters μ_μ and σ_μ, we use the same informal empirical Bayes strategy in CGM10. Based on the idea that that $E(Y \,|\, x)$ is very likely between y_{min} and y_{max}, the observed minimum and maximum of Y, we want to choose μ_μ and σ_μ so that the induced prior on $E(Y \,|\, x)$ assigns substantial probability to the interval (y_{min}, y_{max}). By using the observed y_{min} and y_{max}, we aim to ensure that the implicit prior for $E(Y \,|\, x)$ is in the right "ballpark".

In the unconstrained case where each value of $E(Y \,|\, x)$ is the sum of m iid μ_{ij}'s under the sum-of-trees model, the induced prior on $E(Y \,|\, x)$ under (3.5) is exactly $N(m\,\mu_\mu, m\,\sigma_\mu^2)$. Let us argue now that when monotone constraints are introduced, $N(m\,\mu_\mu, m\,\sigma_\mu^2)$ still holds up as a useful approximation to the induced prior on $E(Y \,|\, x)$. To begin with, for each value of x, let $g(x; T_j, M_j) = \mu_{xj}$, the mean assigned to x by the jth tree T_j. Then, under the sum-of-trees model, $E(Y \,|\, x) = \sum_{j=1}^m \mu_{xj}$ is the sum of m independent means since the μ_{xj}'s

are independent across trees. Using central limit theorem considerations, this sum of small random effects will be approximately normal, at least for the central part of the distribution. The means of all the random effects will be centered around μ_μ, (the constrained μ_{ij}'s will have pairwise offsetting biases), and so the mean of $E(Y \mid x)$ will be approximately μ_μ. Finally, since the marginal variance for all μ_{xj}'s is at least approximately σ_μ^2, the variance of $E(Y \mid x)$ will be approximately $m\sigma_\mu^2$.

Proceeding as in CGM10, we thus choose μ_μ and σ_μ so that $m\,\mu_\mu - k\sqrt{m}\,\sigma_\mu = y_{min}$ and $m\,\mu_\mu + k\sqrt{m}\,\sigma_\mu = y_{max}$ for some preselected value of k. This is conveniently implemented by first shifting and rescaling Y so that the observed transformed y values range from $y_{min} = -0.5$ to $y_{max} = 0.5$, and then setting $\mu_\mu = 0$ and $\sigma_\mu = 0.5/k\sqrt{m}$. Using $k = 2$, for example, would yield a 95% prior probability that $E(Y \mid x)$ is in the interval (y_{min}, y_{max}), thereby assigning substantial probability to the entire region of plausible values of $E(Y \mid x)$ while avoiding overconcentration and overdispersion. As k and/or the number of trees m is increased, this prior will become tighter, thus limiting the effect of the individual tree components of (2.1) by keeping the μ_{ij} values small. We have found that values of k between 1 and 3 yield good results, and we recommend $k = 2$ as an automatic default choice, the same default recommendation for BART. Alternatively, the value of k may be chosen by cross-validation from a range of reasonable choices.

3.3.1 The Choice of m

Again as in BART, we treat m as a fixed tuning constant to be chosen by the user. For prediction, we have found that mBART performs well with values of at least $m = 50$. For variable selection, values as small as $m = 10$ are often effective.

4 Simulation from the Constrained Posterior

4.1 Backfitting MCMC for Constrained Regression Trees

Let y be the $n \times 1$ vector of independent observations of Y from (2.1). All post-data information for Bayesian inference about any aspects of the unknowns, $(T_1, M_1), \ldots, (T_m, M_m)$, σ and future values of Y, is captured by the full posterior distribution

$$p((T_1, M_1), \ldots, (T_m, M_m), \sigma \mid y). \tag{4.1}$$

Since all inference is conditional on the given predictor values, we suppress them in the notation. This posterior is proportional to the product of the likelihood $p(y\,|(T_1, M_1), \ldots, (T_m, M_m), \sigma)$, which is the product of normal likelihoods based on (2.1), and the constrained regularization prior $p((T_1, M_1), \ldots, (T_m, M_m), \sigma)$ described in Section 3.

To extract information from (4.1), which is generally intractable, we propose an MCMC backfitting algorithm that simulates a sequence of draws, $k = 1, \ldots, K$,

$$(T_1, M_1)^{(k)}, \ldots, (T_m, M_m)^{(k)}, \sigma^{(k)} \tag{4.2}$$

that is converging in distribution to (4.1) as $K \to \infty$.

Beginning with a set of initial values of $((T_1, M_1)^{(0)}, \ldots, (T_m, M_m)^{(0)}, \sigma^{(0)})$, this algorithm proceeds by simulating a sequence of transitions $(T_j, M_j)^{(k)} \to (T_j, M_j)^{(k+1)}$, for $j = 1, \ldots, m$, $\sigma^{(k)} \to \sigma^{(k+1)}$. The $(T_j, M_j)^{(k)} \to (T_j, M_j)^{(k+1)}$ transition is obtained by using a Metropolis-Hastings (MH) algorithm to simulate a single transition of a Markov chain with stable distribution

$$p((T_j, M_j)\,|\,r_j^{(k)}, \sigma^{(k)}), \tag{4.3}$$

for $j = 1, \ldots, m$, where

$$r_j^{(k)} \equiv y - \sum_{j'<j} g(x; T_{j'}, M_{j'})^{(k+1)} - \sum_{j'>j} g(x; T_{j'}, M_{j'})^{(k)} \tag{4.4}$$

is the $n-$vector of partial residuals based on a fit that excludes the most current simulated values of $T_{j'}, M_{j'}$ for $j' \neq j$. A full iteration of the algorithm is then completed by simulating the draw of $\sigma^{(k+1)}$ from the full conditional

$$\sigma\,|\,(T_1, M_1)^{(k+1)}, \ldots, (T_m, M_m)^{(k+1)}, y. \tag{4.5}$$

Because conditioning the distribution of (T_j, M_j) on $r_j^{(k)}$ and $\sigma^{(k)}$ in (4.3) is equivalent to conditioning on the excluded values of $(T_{j'}, M_{j'})$, $\sigma^{(k)}$ and y, this algorithm is an instance of MH within a Gibbs sampler.

4.2 A New Localized MH Algorithm

We now describe a new localized MH algorithm for the simulation of $(T_j, M_j)^{(k)} \to (T_j, M_j)^{(k+1)}$ as single transitions of a Markov chain converging to the (possibly constrained) posterior (4.3). For simplicity of notation, let us denote a generic instance of these moves by (T^0, M^0) $\to (T^1, M^1)$. Dropping $\sigma^{(k)}$ from (4.3) since it is fixed throughout this move, and dropping all

the remaining subscripts and superscripts, the target posterior distribution can be expressed as

$$p(T, M \mid r) = p(r \mid T, M)p(M \mid T)p(T)/p(r), \tag{4.6}$$

where its components are as follows.

First, $p(r \mid T, M)$ is the normal likelihood which would correspond to an observation of $r = g(x; T, M) + \epsilon$, where $\epsilon \sim N_n(0, \sigma^2 I)$. Assuming $M = (\mu_1, \ldots, \mu_b)$, and letting r_i be the vector of components of r assigned to μ_i by T, this likelihood is of the form

$$p(r \mid T, M) = \prod_{i=1}^{b} p(r_i \mid \mu_i) \tag{4.7}$$

where

$$p(r_i \mid \mu_i) \propto \prod_j \exp(-(r_{ij} - \mu_i)^2/2\sigma^2). \tag{4.8}$$

The prior of $M \mid T$ given by (3.3) is of the form

$$p(M \mid T) \propto \left[\prod_{i=1}^{b} p(\mu_i \mid T)\right] \chi_C(T, M), \tag{4.9}$$

where $p(\mu_i \mid T) = \phi_{\mu_\mu, \sigma_\mu}(\mu_i)$ from (3.5) if μ_i is unconstrained by χ_C, and $p(\mu_i \mid T) = \phi_{\mu_\mu, c\sigma_\mu}(\mu_i)$ from (3.6) if μ_i is constrained by χ_C. The tree prior $p(T)$ described in Section 3.1 is the same form used for unconstrained BART. Finally, the intractable marginal $p(r)$, which would in principle be obtained by summing and integrating over T and M, will fortunately play no role in our algorithm.

In unconstrained CART and BART, CGM98 and CGM10 used the following two step Metropolis-Hastings (MH) procedure for the simulation of $(T^0, M^0) \to (T^1, M^1)$. First, a proposal T^* was generated with probability $q(T^0 \to T^*)$. Letting $q(T^* \to T^0)$ be the probability of the reversed step, the move $T^1 = T^*$ was then accepted with probability

$$\begin{aligned} \alpha &= \min\left\{\frac{q(T^* \to T^0)}{q(T^0 \to T^*)} \frac{p(T^* \mid r)}{p(T^0 \mid r)},\ 1\right\} \\ &= \min\left\{\frac{q(T^* \to T^0)}{q(T^0 \to T^*)} \frac{p(r \mid T^*)}{p(r \mid T^0)} \frac{p(T^*)}{p(T^0)},\ 1\right\}. \end{aligned} \tag{4.10}$$

If accepted, any part of M^1 with a new ancestry under M^1 is simulated from independent normals since $p(M \mid T^1, r)$ just consists of b independent normals given the independence and conditional conjugacy of our prior (which is (4.9) without the monotonicity constraint

$\chi_C(T, M)$) and the conditional data independence (4.7). Otherwise (T^1, M^1) is set equal to (T^0, M^0).

In the contrained case, the basic algorithm is the same except that with the monotonicity constraint in (4.9), the μ_i in M are dependent. Hence, when we make local moves involving a few of the μ_i we must be careful to condition on the remaining elements. In addition, computations must be done numerically since we lose the conditional conjugacy. The moves in mBART only operate on one or two of the μ values at a time so that the appropriate conditional integrals can easily be done numerically.

We consider localized proposals $(T^0, M^0) \to (T^*, M^*)$ under which M^0 and M^* differ only by those μ's which have different ancestries under T^0 and T^*. Letting μ_{same} be the part of M^0 with the same ancestry under T^0 and T^*, we restrict attention to proposals for which $M^0 = (\mu_{same}, \mu_{old})$ and $M^* = (\mu_{same}, \mu_{new})$, where μ_{old} is the part of M^0 that will be replaced by μ_{new} in M^*. It will also be convenient in what follows to let r_{old} be the components of the data r assigned to μ_{old} by T^0, r_{new} to be the components assigned to μ_{new} by T^*, and r_{same} to be the components assigned to the identical components of μ_{same} by both T^0 and T^*.

For example, suppose we begin with a proposal $T^0 \to T^*$ that randomly chooses between a birth step and death step, and that T^* was obtained by a birth step, which entails adding two child nodes at a randomly chosen terminal node of T^0. This move is illustrated in Figure 4 where $M^0 = (\mu_1, \mu_2, \mu_0)$ and $M^* = (\mu_1, \mu_2, \mu_L, \mu_R)$, so that $\mu_{same} = (\mu_1, \mu_2)$ to which $r_{same} = (r_1, r_2)$ is assigned, $\mu_{old} = \mu_0$ to which $r_{old} = r_0$ is assigned, and $\mu_{new} = (\mu_L, \mu_R)$ to which $r_{new} = (r_L, r_R)$ is assigned. Note that the set of observations in (r_L, r_R) is just the division of the set of observations in r_0 defined by the decision rule associated with node 7 in the tree T^*.

The key is to then proceed conditional on μ_{same} and the tree ancestry associated with it. In Figure 4, we condition on $\mu_{same} = (\mu_1, \mu_2)$ *and* the ancestral tree structure given by nodes $(1, 2, 3, 6)$ including the decision rules associated with the interior nodes 1 and 3. To keep the notation clean, we will use μ_{same} as a conditioning variable in our expressions below and the reader must make a mental note to include the associated tree ancestry as conditioning information.

Conditional on μ_{same}, our Metropolis procedure is as follows. First, a proposal T^* is generated with probability $q(T^0 \to T^*)$, using the same CGM98 proposal used in unconstrained CART and BART. Letting $q(T^* \to T^0)$ be the probability of the reversed step, the move $T^1 = T^*$

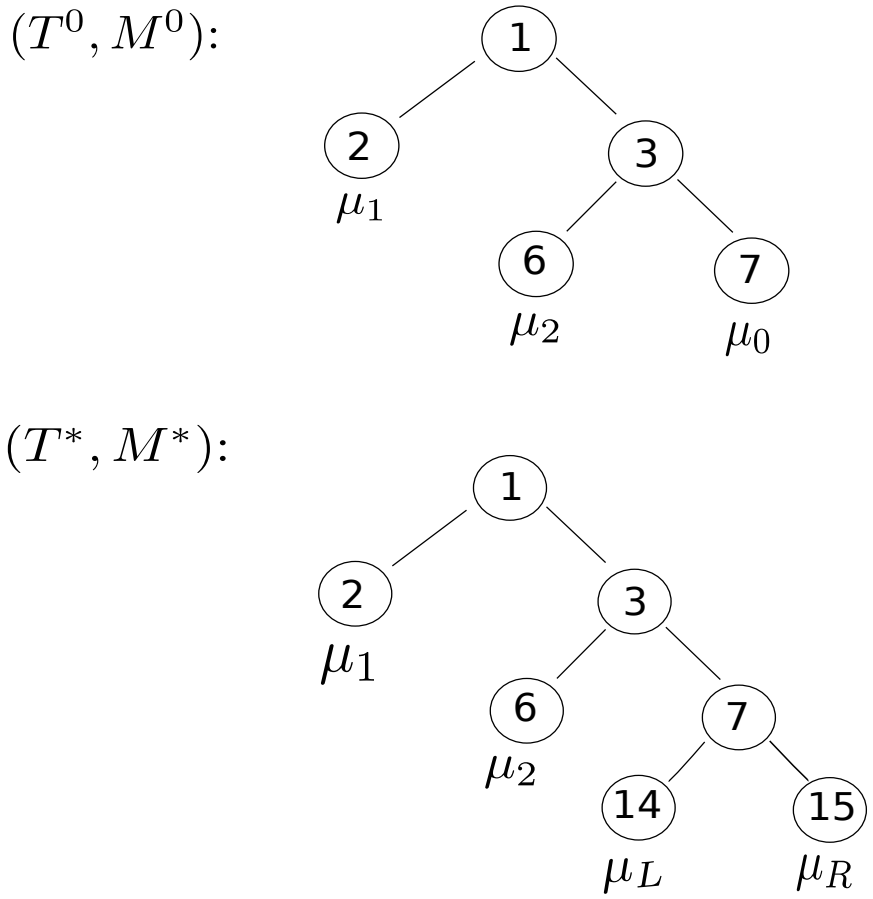

Figure 4: A typical birth step starting at (T^0, M^0) and proposing (T^*, M^*). T^0 includes the nodes 1,2,3,6,7. T^* includes the nodes 1,2,3,6,7,14,15. Here $\mu_{same} = (\mu_1, \mu_2)$. Our MH step proceeds conditionally on μ_{same} and the associated ancestral parts of the tree structures T^0 and T^*, nodes 1,2,3,6. Our proposal generates the candidate rule associated with node 7 in T^*. Conditional on all these elements, we integrate out μ_0 or (μ_L, μ_R) subject to the constraints implied by the conditioning elements. Note that the proposal for the node 7 rule does not depend on μ_{same}, it only depends on the tree structures.

is then accepted with probability

$$\begin{aligned}
\alpha &= \min\left\{\frac{q(T^* \to T^0)}{q(T^0 \to T^*)}\,\frac{p(T^* \mid \mu_{same}, r)}{p(T^0 \mid \mu_{same}, r)},\ 1\right\} \\
&= \min\left\{\frac{q(T^* \to T^0)}{q(T^0 \to T^*)}\,\frac{p(T^* \mid \mu_{same}, r_{new})}{p(T^0 \mid \mu_{same}, r_{old})},\ 1\right\} \\
&= \min\left\{\frac{q(T^* \to T^0)}{q(T^0 \to T^*)}\,\frac{p(r_{new} \mid T^*, \mu_{same})}{p(r_{old} \mid T^0, \mu_{same})}\frac{p(T^*)}{p(T^0)},\ 1\right\}. \qquad (4.11)
\end{aligned}$$

The difference between (4.10) and (4.11) is that we condition on μ_{same} throughout and explicitly note that the r_{same} part of r does not matter. In going from the first line above to the second we have used the fact that, conditional on μ_{same}, r_{same} gives the same multiplicative contribution to the top and bottom of the acceptance ratio so that it cancels out leaving only terms depending on r_{new} and r_{old}. To go from the second line above to the third we will compute the required r_{new} and r_{old} marginals numerically as detailed in Section 4.3 below. Note also that in the BART prior, T and M are dependent only through the dimension of M so $p(T^*)$ / $p(T^0)$ is the same as in the unconstrained case.

If $T^1 = T^*$ is accepted, μ_{new} is then simulated from $p(\mu_{new} \,|T^1, \mu_{same}, r) = p(\mu_{new} \,|T^1, \mu_{same}, r_{new})$ and M^1 is set equal to (μ_{same}, μ_{new}). Otherwise (T^1, M^1) is set equal to (T^0, M^0).

4.3 Implementation of the Localized MH Algorithm

The implementation of our localized MH algorithm requires the evaluation of $p(r_{new} \,|T^*, \mu_{same})$ and $p(r_{old} \,|T^0, \mu_{same})$ for the α calculation in (4.11), and the simulation from $p(\mu_{new} \,|T^1, \mu_{same}, r_{new})$. Although these can all be done quickly and easily in the unconstrained cases, a different approach is needed for constrained cases. This approach, which we now describe, relies crucially on the reduced computational requirements for the localized MH algorithm when $T^0 \to T^*$ is restricted to local moves at a single node.

For the moment, consider the birth move described in Section 4.2 and illustrated in Figure 4. In this case, $\mu_{new} = (\mu_L, \mu_R)$ with corresponding $r_{new} = (r_L, r_R)$ and $\mu_{old} = \mu_0$ with corresponding r_0. Thus, to perform this move, it is necessary to compute $p(r_L, r_R \,|\, T^*, \mu_{same})$ and $p(r_0 \,|\, T^0, \mu_{same})$ for the computation of α in (4.11), and to simulate (μ_L, μ_R) from $p(\mu_L, \mu_R \,|\, r_L, r_R, T^*, \mu_{same})$ when $T^1 = T^*$ is selected. For the corresponding death step, we would need to simulate μ_0 from $p(\mu_0 \,|\, r_0, T^0, \mu_{same})$. When these means are unconstrained, these calculations can be done quickly with closed form expressions and the simulations by routine methods so we focus here on the constrained case.

Let us begin with the calculation of

$$p(r_L, r_R \,|\, T^*, \mu_{same}) = \int p(r_L \,|\, \mu_L)\, p(r_R \,|\, \mu_R)\, p(\mu_L, \mu_R \,|\, T^*, \mu_{same})\, d\mu_L\, d\mu_R \tag{4.12}$$

where

$$p(\mu_L, \mu_R \,|\, T^*, \mu_{same}) = \phi_{\mu_\mu, c\sigma_\mu}(\mu_L)\, \phi_{\mu_\mu, c\sigma_\mu}(\mu_R)\, \chi_C(\mu_L, \mu_R) \;/\; d_* \tag{4.13}$$

and d_* is the normalizing constant. The determination of $\chi_C(\mu_L, \mu_R)$ is discussed in Section 2; it is the set $(\mu_L, \mu_R, \mu_{same})$ which results in a monotonic function. Note that C is of the form $C = \{(\mu_L, \mu_R) : a \le \mu_L \le \mu_R \le b\}$ with a, b (possibly $-\infty$ and/or ∞) determined by the conditioning on T^* and μ_{same}. In particular, note that C depends on μ_{same} but we have suppressed this in the notation for the sake of simplicity.

Closed forms for (4.12) and the norming constant d_* are unavailable. However, since the integrals are only two-dimensional, it is straighforward to compute them numerically. To use a very simple approach, we approximate them by summing over a grid of (μ_L, μ_R) values. We choose a grid of equally spaced μ values and then let G be the set of (μ_L, μ_R) where both μ_L and μ_R belong to the grid. Then, our approximate integrals are

$$\tilde{p}(r_L, r_R \,|\, T^*, \mu_{same}) = \sum_{(\mu_L, \mu_R) \in G \cap C} p(r_L \,|\, \mu_L)\, p(r_R \,|\, \mu_R)\, \tilde{p}(\mu_L, \mu_R \,|\, T^*, \mu_{same}), \tag{4.14}$$

where

$$\tilde{p}(\mu_L, \mu_R \,|\, T^*, \mu_{same}) = \phi_{\mu_\mu, c\sigma_\mu}(\mu_L) \phi_{\mu_\mu, c\sigma_\mu}(\mu_R) \;/\; \tilde{d}_* \tag{4.15}$$

with

$$\tilde{d}_* = \sum_{(\mu_L, \mu_R) \in G \cap C} \phi_{\mu_\mu, c\sigma_\mu}(\mu_L) \phi_{\mu_\mu, c\sigma_\mu}(\mu_R). \tag{4.16}$$

Note that we do not include "$\Delta\mu$" terms (the difference between adjacent grid values) in our integral approximations since they cancel out.

If $T^1 = T^*$ is accepted, the simulation of (μ_L, μ_R) proceeds by sampling from the probability distribution over $G \cap C$ given by

$$\tilde{p}(\mu_L, \mu_R \,|\, r_L, r_R, T^*, \mu_{same}) = \frac{p(r_L \,|\, \mu_L)\, p(r_R \,|\, \mu_R)\, \tilde{p}(\mu_L, \mu_R \,|\, T^*, \mu_{same})}{\tilde{p}(r_L, r_R \,|\, T^*, \mu_{same})}. \tag{4.17}$$

Note that $\tilde{d}_*$ cancels in (4.17) so that we are just renormalizing

$$p(r_L \,|\, \mu_L)\, p(r_R \,|\, \mu_R)\, \phi_{\mu_\mu, c\sigma_\mu}(\mu_L) \phi_{\mu_\mu, c\sigma_\mu}(\mu_R)$$

to sum to one on $G \cap C$.

For the calculation of

$$p(r_0 \mid T^0, \mu_{same}) = \int p(r_0 \mid \mu_0)\, p(\mu_0 \mid T^0, \mu_{same})\, d\mu_0 \tag{4.18}$$

where

$$p(\mu_0 \mid T^0, \mu_{same}) = \phi_{\mu_\mu, c\sigma_\mu}(\mu_0)\chi_C(\mu_0) \;/\; d_0 \tag{4.19}$$

and d_0 is the normalizing constant with the constraint set of the form $C = \{(\mu_0) : a \leq \mu_0 \leq b\}$, similar griding can be done to obtain a discrete approximation $\tilde{d}_0$ of d_0 and a constrained posterior sample of μ_0. Again, C implicitly depends on T^0 and μ_{same}. The grid here would be just one-dimensional.

Computations for the reverse death move would proceed similarly. Local moves for $T^0 \to T^*$ beyond birth and death moves may also be similarly applied, as long as μ_{old} and μ_{new} are each at most two dimensional since beyond two dimensions, grids become computationally demanding. For example, $T^0 \to T^*$ obtained by changing a splitting rule whose children are terminal nodes would fall into this category. In all our examples, we use birth/death moves and draws of a single μ component given T and all the remaining elements of M.

The approach outlined above for birth/death moves involves two bivariate integrals and two univariate integrals which we approximate with two sums over a bivariate grid and two sums over a univariate grid. In practice, we reduce the computational burden by letting $\tilde{d}_*$ and $\tilde{d}_0$ equal one and then compensating for this omission with an adjustment of our T prior. For example, in a birth move, setting the d's to one ignores a factor $\tilde{d}_0 \;/\; \tilde{d}_*$ in our ratio. Note that from (4.13) d^* is just the constrained integral of the product of two univariate normal densities. Without the constraint, the integral would be one. The more our monotonicity constraint limits the integral (through $\chi_C(\mu_L, \mu_R)$), the smaller d^* is. Similary, d_0 is a constrained univariate integral. However, in a birth step, d_* is typically more constrained than d_0. Hence, $\tilde{d}_0 \;/\; \tilde{d}_*$ is a ratio depending on T^0 and T^* which we expect to be great than one. Note that d_0 only depends on T^0 and d_* only depends on T^* (that is, not on μ_{same}).

We compensate for the omission of $\tilde{d}_*$ and $\tilde{d}_0$ by letting $\alpha = .25$ and $\beta = .8$ rather than using standard BART default values of $\alpha = .95$ and $\beta = 2$. With $\alpha = .25$ and $\beta = .8$, $p(T^*)/p(T^0)$ is larger mimicking the effect of the omitted d ratio. We have found that with these choices we get tree sizes comparable to those obtained in unconstrained BART. The values $\alpha = .25$ and $\beta = .8$ are used in all our examples.

5 Examples

In this section we present several examples to illustrate the performance of mBART. We present results for two simulated scenarios and three real data sets. In the simulations, where we know the true function is monotonic and non-linear, we compare mBART to BART. For the real data, we compare mBART to BART and the standard linear model. In all cases (except Figure 6 in which we study prior sensitivity) we use default priors for mBART and BART but remind the reader that for best out-of-sample results, it may be wise to consider the use of cross-validation to tune the prior choice as illustrated in CGM10.

Note that below when we refer to the "fit" of BART or mBART at a given x, we mean the posterior mean of $f(x)$ estimated by averaging the f draws evaluated at x.

Our first simulated example has a one-dimensional x so that we can easily see the fits graphically. We see three differences between mBART and BART which are intuitive consequences of the the injection of the strong prior information that the function is monotonic: (i) the fitted function is smoother, (ii) the uncertainty is lower, and (iii) for some values of x, the influence of the prior is reduced.

Our second example simulates data sets where f is a five-dimensional nonlinear monotonic function. We explore the relationship between the out-of-sample predictive performance of mBART and the signal to noise ratio by using four different values for the error standard deviation. In the high signal case, BART is able to estimate f without the aid of the additional information about monotonicity so that BART and mBART have the same performance out-of-sample. In the low signal cases, the additional information is important and mBART beats BART out-of-sample.

In our first real example y is the price of a house and x represents attributes of the house. In this example, mBART, BART, and the linear model all give similar fits. However, the non-monotonic BART fits are very counter intuitive while the mBART fits are quite reasonable. This example also illustrates a very important feature of mBART: it handles *ordered* categorical explanatory variables beautifully.

In our second real example y is the price of a used car and x represents attributes of a car. In this example, the linear model fails while mBART and BART perform comparably with mBART giving more interpretable results.

In our third real example, y is the excess return for a cross section of firms in a given month and x measures attributes of each firm in the previous month. This example is extreme in

that the signal to noise ratio is very low. The unconstrained BART fit is very noisy while the mBART fit is nicely smooth and suggestive of non-linearity.

The examples illustrate two key attributes of mBART in cases where the monotonicity is appropriate. First, the smoother mBART fits are more interpretable than the more "jumbly" BART fits. Second, in terms of the bias-variance trade off, the monotone constraint may serve to restrain the fit giving improved out-of-sample performance. This is clearly illustrated in our second simulated example. In all three of our real examples, the posterior distribution of σ from BART covers the posterior mean from mBART, informally suggesting that the monotonicity constraint is consistent with the data. In all three of our real examples, the mBART fits are more appealing. In the car price example, BART and mBART beat the linear model out-of-sample, and mBART is (perhaps) a bit better than BART. In the returns example, linear beats BART out-of-sample, while mBART does as well as linear. In general, we can think of mBART as a convenient "half way house" in between the very flexible ensemble method BART and the inflexible linear method.

5.1 One-Dimensional Simulated Example

In this section we present a very simple simulated example with just one x variable so that we can visually see some of the basic properties of the monotone BART inference. We simulate $n = 200$ observations from the model

$$Y_i = X_i^3 + \epsilon_i, \quad \epsilon_i \sim N(0, .1^2),\ X_i \sim U[-1, 1].$$

Figure 5 shows the results for a single simulated data set. The BART inference is displayed in the left panel and the mBART inference is in the right panel. The mBART fit is much better. The fit is smoother and the uncertainty is much smaller. Clearly, injecting the correct prior information that the function is monotone dramatically tightens up the inference.

Given the additional information in the monotonicity constraint, mBART should also be less sensitive to the choice of prior than BART. To explore this, we fit BART and mBART using 36 different prior specifications and compared the variation in fits. We used three different values for m, the number of trees in the sum: 50, 200, and 500. We used four different values for k where a large value of k implies a tighter prior on the end node μ parameters : $k = 1$ 1, 2, 3, or 5 (see CGM10 and Section (3.3)). We used three different settings for the pair of values (ν, q) where ν is the degrees of freedom for the inverted chi-squared prior and the scale is chosen so the least squares estimate of the error standard deviation σ is at the q^{th}

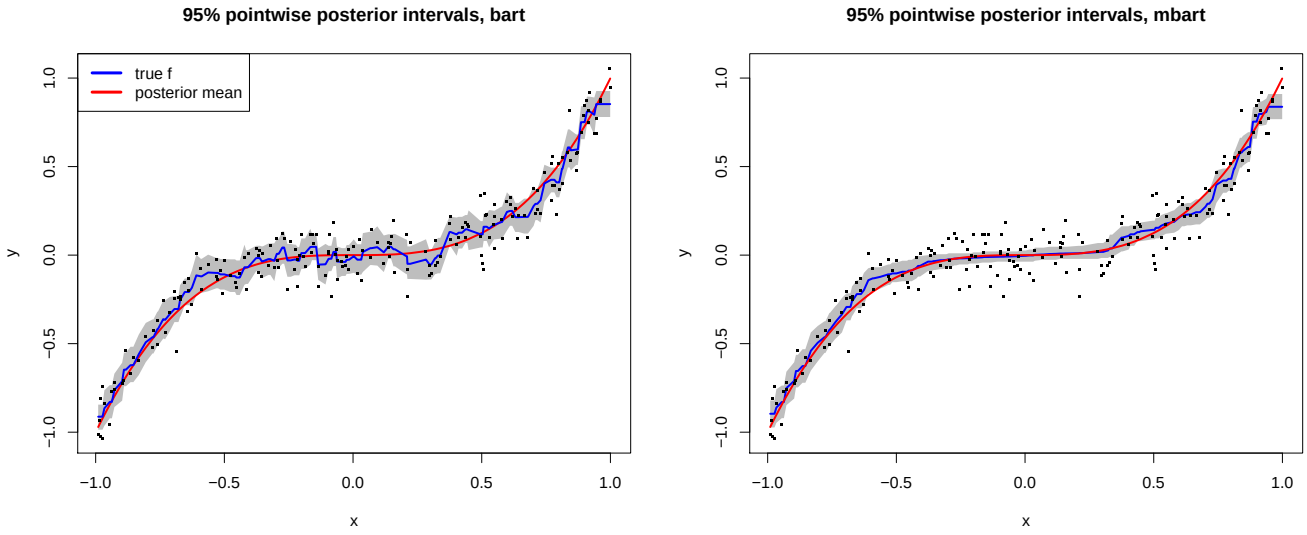

Figure 5: *One-dimensional simlulation. The left panel shows the unconstrained BART fit. The right panel shows the mBART fit. The monotone fit is much better and the 95% pointwise intervals for $f(x)$ are much tighter.*

quantile of the σ prior: (3,.9), (3,.99), and (10,.75). These settings follow the prior settings explored in CGM10. All possible combinations gives us $3 \times 4 \times 3 = 36$ possible settings.

The results for one simulated data set are shown in Figure 6. The horizontal axis is observation number which goes from 1 to 200. Before we ran the methods we sorted the observations according to the value of x. Hence, x is increasing as you go from left to right in Figure 6. For each observation the left vertical line gives the range of fitted values for BART while the right vertical line gives the range of fitted values for mBART. The average fitted value is subtracted from the range so we can see the relative spreads.

Interestingly, which method gives a greater dispersion depends on the location of the x at which predictions are made. In the middle of the data mBART is dramatically less sensitive to the prior choice. However at the end points, mBART is more sensitive. This make sense as the monotonicity constrains points in the interior but not at the edges.

5.2 Five-Dimensional Simulated Example

In this section we present results for a five-dimensional simulated example. We compare the out-of-sample prediction for mBART with that of BART. We vary the size of the error

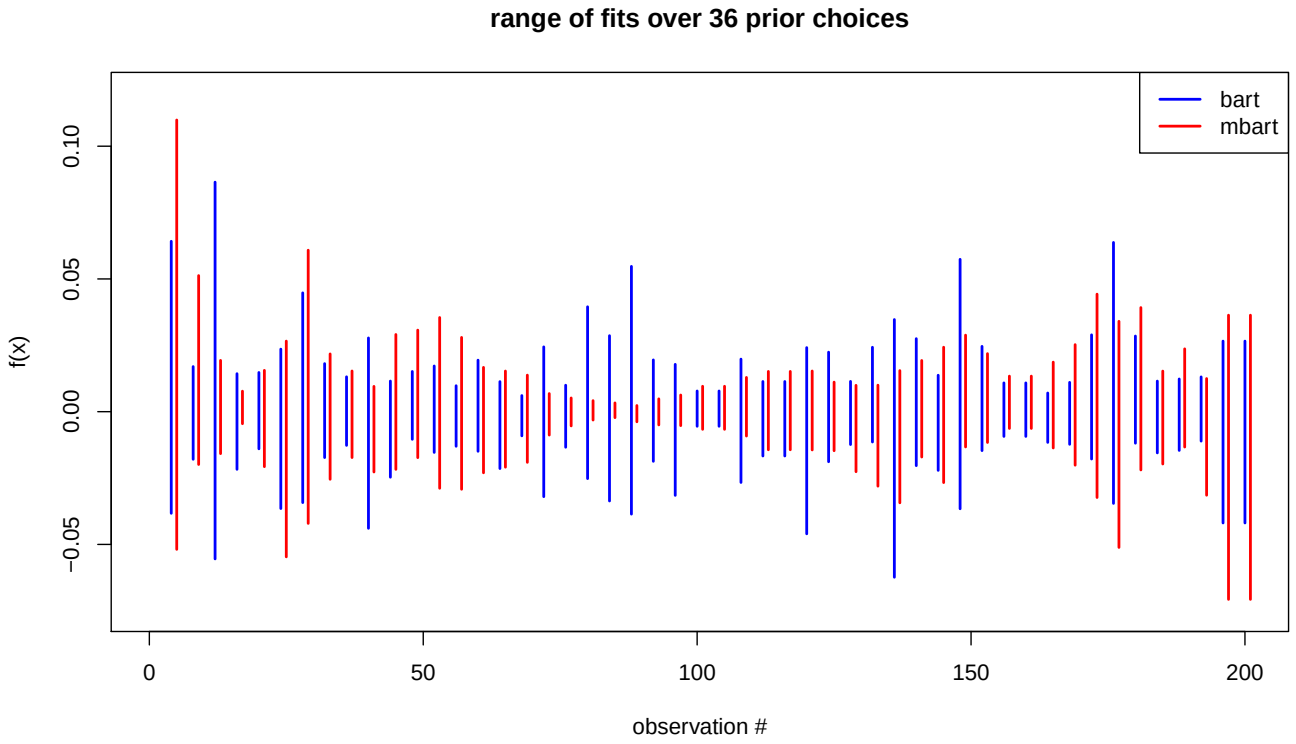

Figure 6: *One-dimensial simulation. Prior sensitivity. The range of fits obtained at each x value using different prior specifications.*

standard deviation in $Y = f(x) + \epsilon$ where x denotes a vector which five components. For small errors, there is no difference in the performance as BART is able to infer the function with very little error. As the the error variance increases, the additional information that the function is monotonic becomes more useful and mBART outperforms BART.

We simulated the data using

$$Y_i = x_{i1}x_{i2}^2 + x_{i3}x_{i4}^3 + x_{i5} + \epsilon_i.$$

All of the x values are drawn from the standard uniform distribution on (0,1). Our function $f(x) = x_1x_2^2 + x_3x_4^3 + x_5$ is clearly monotonic for such x. We let $\epsilon_i \sim N(0, \sigma^2)$ with $\sigma = .2$, .5, .7, and 1. For each value of σ we simulated 200 data sets. Each data set has 500 in-sample (training) observations and 1,000 out-of-sample (test) observations. For the training data, we draw x and y, while for the test data we need only draw x. For each simulated data set we compute the RMSE

$$RMSE = \sqrt{\frac{1}{1000}\sum_{i=1}^{1000}(f(x_i) - \hat{f}(x_i))^2}$$

where f is the true function, $\hat{f}(x_i)$ is the posterior mean, and the x_i are the test x vectors.

The results are displayed in Figure 7. Each boxplot depicts the 200 RMSE values for a method. The first two boxplots show results for the draws with $\sigma = .2$, the second two for $\sigma = .5$ and so on. Within each pair of boxplots, the first boxplot gives the unconstrained BART RMSE's and the second one gives the mBART RMSE's. For the smallest value of σ both methods give similar results. As σ increases, mBART clearly outperforms unconstrained BART.

5.3 The House Price Example

We have 128 observations on houses that have sold. The dependent variable y is the sales price (thousands of dollars). The explanatory variables in x tell us about the houses. The first explanatory variable (`nbhd`) indicates which of three neighborhoods the house is in. The three neighborhoods are labeled 1,2, and 3 and *we know* that neighborhood 1 is more desirable than 2 and 2 than 3 (*location, location, location !!*). The second x is the size of the house (`size`, thousands of square feet). The third x is a 0-1 dummy variable indicating whether or not the house is made of brick (`brickYes`, 0 for non-brick, 1 for brick), and the

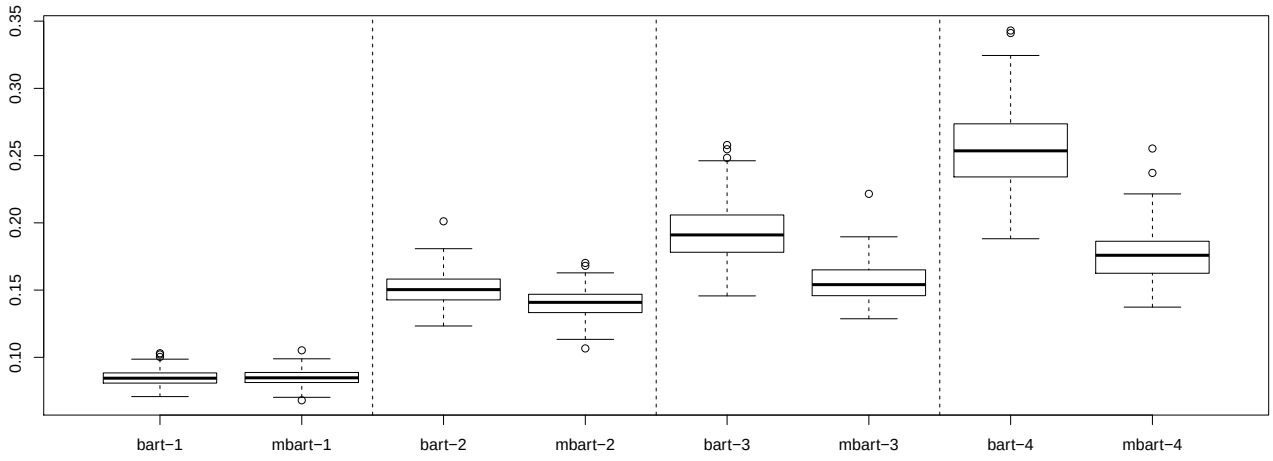

Figure 7: *Five-dimensional simulation. Out-of-sample RMSE.*

fourth x is the number of bedrooms `nbedrm`. In each case, a monotone relationship seems likely with larger x values giving larger y values.

Note that because mBART is based on trees and `nbhd` is an *ordered* categorical variable, we can just include it in mBART as a numeric variable. Of course, this might not be appropriate in a linear model since it would assume the difference between neighborhoods 1 and 2 is the same as the difference between neighborhoods 2 and 3.

Below is the multiple regression output where dummies for neighborhoods 2 and 3 have been included.

```
            Estimate Std. Error t value Pr(>|t|)
(Intercept)   18.130     10.732   1.689   0.0937 .
nbhd2          5.096      2.788   1.828   0.0700 .
nbhd3         34.909      3.232  10.802  < 2e-16 ***
size          42.656      6.031   7.073 1.03e-10 ***
brickYes      19.363      2.433   7.958 1.02e-12 ***
nbedrm         2.700      1.925   1.403   0.1632
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

Residual standard error: 12.45 on 122 degrees of freedom
Multiple R-squared:  0.7936,Adjusted R-squared:  0.7851
F-statistic: 93.81 on 5 and 122 DF,  p-value: < 2.2e-16
```

Note that the signs of the coefficients are all positive. However, the relationship need not be linear and we can check this by running the flexible fitters BART and mBART.

Figure 8 displays aspects of the inference from the linear model, BART, and mBART. The top left panel plots the BART MCMC draws of σ. The solid horizontal line indicates the σ estimate from the linear regression, the "dot dash" horizontal line indicates the posterior mean of σ from mBART, and the "dash" horizontal line indicates the posterior mean of σ from BART. The top right panel shows the σ draws from the mBART MCMC with the horizontal lines indicating estimates of σ as in the top left panel. The bottom left panel plots the BART fits versus the mBART fits. The bottom right panel plots the linear fits versus the mBART fits. All three methods give a similar level of fit.

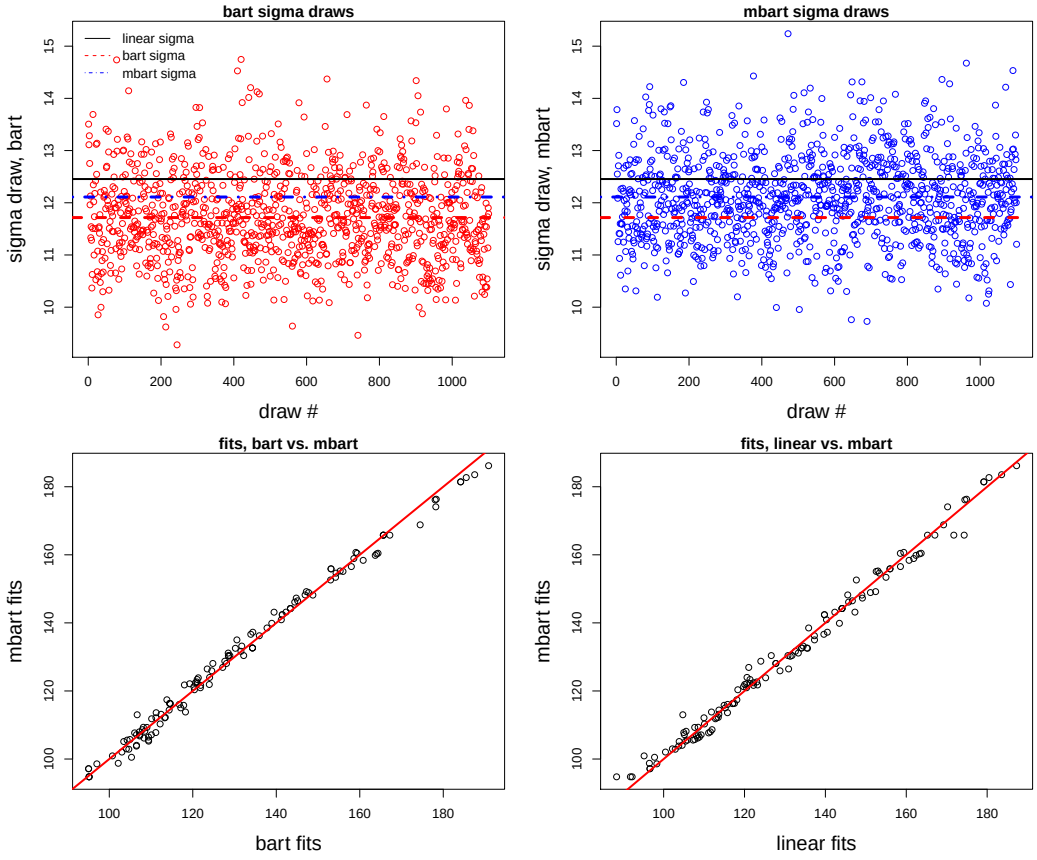

Figure 8: *House price example. Top row: σ draws from BART (left panel), σ draws from mBART (right panel). Solid horizontal line at least squares estimate of σ. Bottom row: BART versus mBART (left panel) and linear fits versus mBART (right panel).*

Figure 9 displays the conditional effects of neighborhood (`nbhd`) and house size (`size`). To

visualize the conditional effects from the BART/mBART fits we construct x vectors such that one of the x coordinates varies while the others are held fixed. For example, in the bottom left panel of Figure 9 we see the estimate of $f(x)$ for x having the size values indicated on the horizontal axis. The various curves in the figure correspond to different fixed levels of the other three variables in x. We picked a grid of values for each variable and then constructed a design matrix composed of all possible combinations. We used 15 quantiles for the values of `size`, 2, 3, 4, and 5 for `nbedrm`, 1,2, and 3 for `nbhd`, and 0 or 1 for `brickYes`. giving $15 \times 4 \times 3 \times 2 = 360$ possible combinations. To keep the plot readable, we randomly sampled the possible combinations of the variables held fixed so that not all possible curves are plotted in each frame of Figure 9.

The BART conditional effects of house size are plotted in the lower left panel of Figure 9 while the mBART estimates are in the lower right panel. We see that the BART estimates suggest that is possible for a larger house to sell for less with all else held constant while the mBART estimates give a much more appealing story. The top two panels of Figure 9 display the neighborhood effect. In this case the data are sufficiently informative for there to be little difference.

Figure 10 displays the conditional effects for brick (top row) and the number of bedrooms (bottom row). The BART estimates are in the left column and the mBART estimates are in the right column. Some BART estimates suggest that price could decrease with additional bedrooms while mBART never does.

The number of bedrooms was entered into the linear regression specification as a numeric variable. This makes the clearly unappealing assumption that going from 2 to 3 bedrooms has the same effect as going from 3 to 4 (at all levels of other variables). We could create dummies with each dummy indicating a particular number of bedrooms, but, in general, this strategy can create a plethora of dummies and we only have 128 observations. In mBART we treat the number of bedrooms as an ordered categorical variable simply by imposing the monotonicty constraint. However, the conclusion is that the number of bedrooms effect is small and that, overall, the linear model may indeed be adequate. In this case mBART serves as a way to validate the linear model. In our next example, the linear model fails.

Fundamental to this paper is the observation that all mBART conditional effects display monotonicity reflecting the model's ability to impose monotonicity in a multivariate setting.

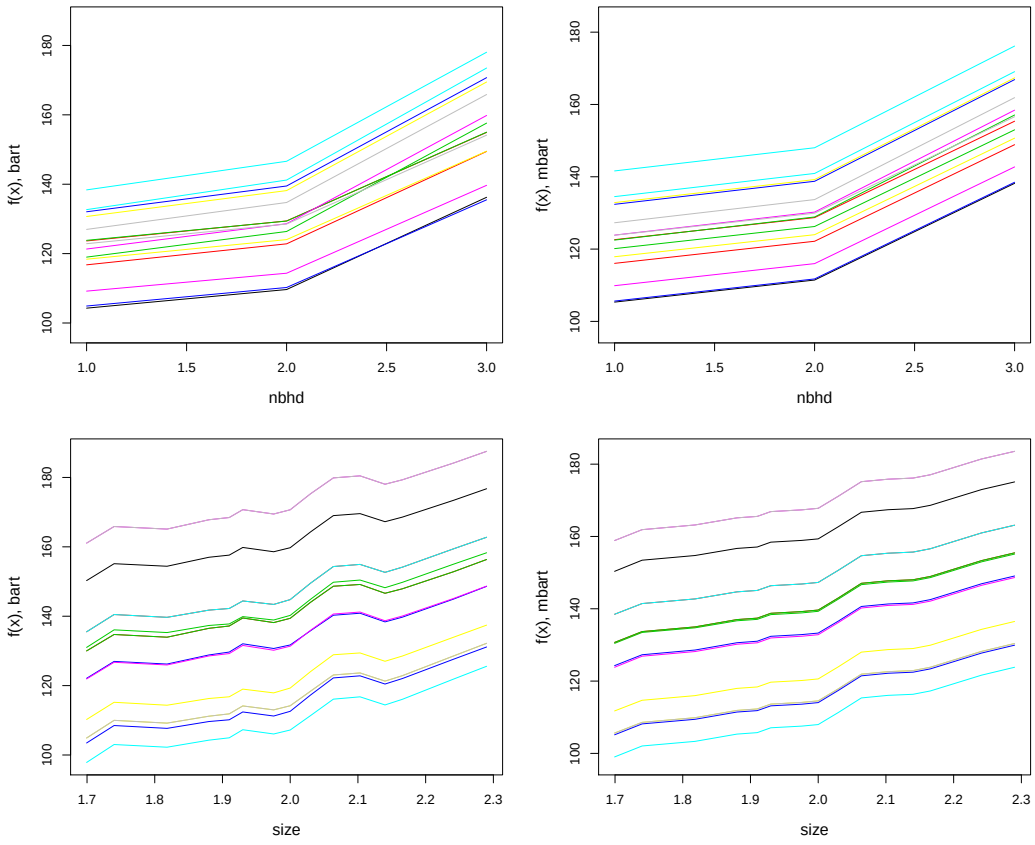

Figure 9: *House price example. Conditional effects of neighborhood (top panels) and house size (bottom panels) for BART (left panels) and mBART (right panels).*

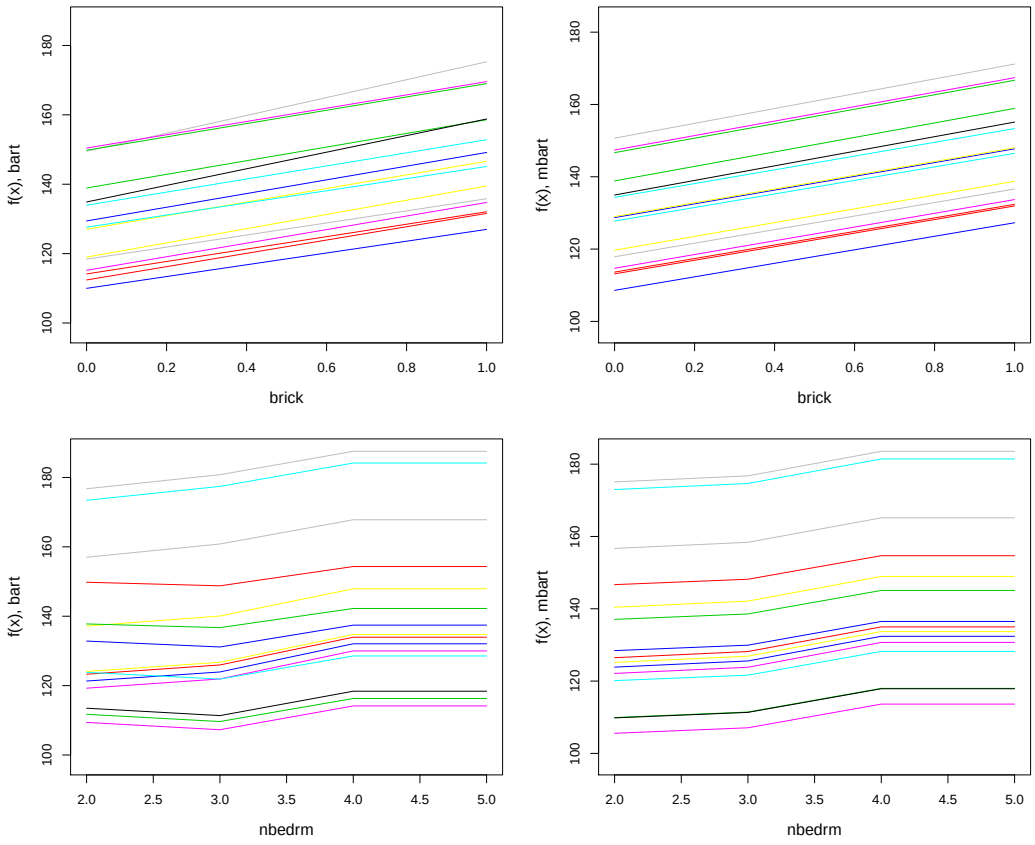

Figure 10: *House price example. Condtional effects of brick (top panels) and number of bedrooms (bottom panels) for BART (left panels) and mBART (right panels).*

5.4 The Car Price Example

In this example we have 1,000 observations and y is the sale price of a used Mercedes car. Our explanatory x variables are: (i) the mileage on the car (`mileage`), (ii) the year of the car (`year`) (iii) feature count (`featureCount`) and (iv) has the car had just one owner (1 if yes, 0 if no) (`isOneOwner`).

We certainly expect a car with more mileage to sell for less. We multiplied mileage by -1 to make the relationship monotonic increasing. We certainly expect the price to be monotonic increasing in `year` and a higher price if there is just one owner. The feature count variable is interesting. Based on our initial understanding of this variable, we expected it to have a negative effect on y so we multiplied by -1 to make it monotonic increasing. Below is the multiple regression output. We see that all the signs are positive and `featureCount` is "significant". It turns out we misunderstood this variable and we will discuss it further when we look at the mBART results. `featureCount` is left in the presented analysis as we feel the case "put in a variable by accident" is very realistic and mBART handles it nicely.

```
Coefficients:
               Estimate Std. Error t value Pr(>|t|)
(Intercept)  -5.427e+06  1.732e+05 -31.334  < 2e-16 ***
mileage       1.529e-01  8.353e-03  18.301  < 2e-16 ***
year          2.726e+03  8.613e+01  31.648  < 2e-16 ***
featureCount  3.263e+01  9.751e+00   3.346 0.000851 ***
isOneOwner    1.324e+03  6.761e+02   1.959 0.050442 .
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

Residual standard error: 7492 on 995 degrees of freedom
Multiple R-squared:  0.8351,Adjusted R-squared:  0.8344
F-statistic:  1260 on 4 and 995 DF,  p-value: < 2.2e-16
```

Figure 11 has the same layout as Figure 8. The top left panel shows σ draws from BART, the top right panel shows σ draws from mBART, and in each plot the estimate of σ from the linear regression is indicated by a horizontal solid line. Both BART and mBART quickly burn-in to σ values much smaller than the least squares estimate indicating a much tighter fit. The monotonicity constraint makes the σ draws slightly larger. In the bottom left panels we see that the BART and mBART fits are quite similar. In the bottom right panel we see that the mBART fits are quite different from the linear model fits.

Figures 12 and 13 have the same layout as Figures 9 and 10. They show the conditional effects for the four variables in x.

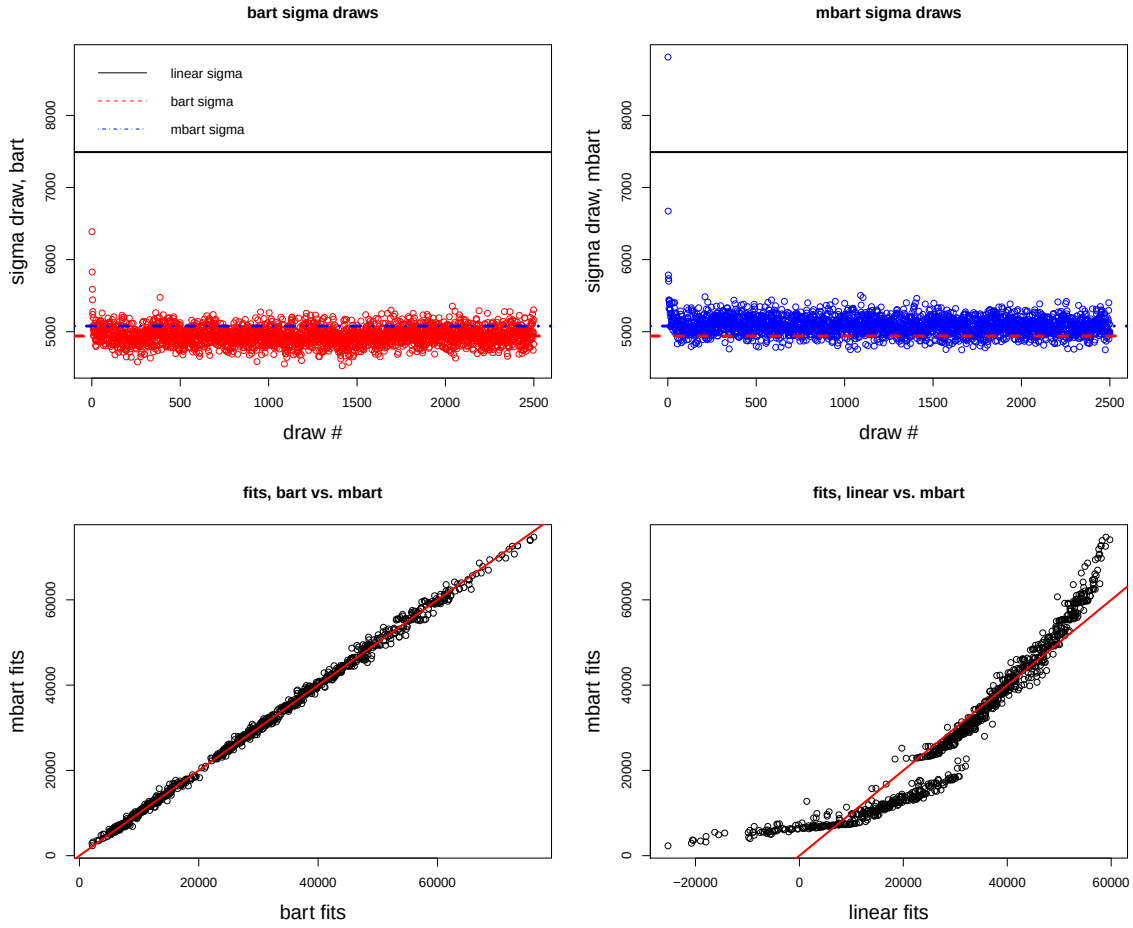

Figure 11: *Car price example. Top row: σ draws from BART (left panel), σ draws from mBART (right panel). Solid horizontal line at least squares estimate of σ. Bottom row: BART versus mBART (left panel) and linear fits versus mBART (right panel).*

The conditional effects for `mileage` and `year` (Figure 12) are quite similar for BART and mBART. The mBART effect for mileage is smoother and everywhere monotonic while the BART fit does exhibit slight dips.

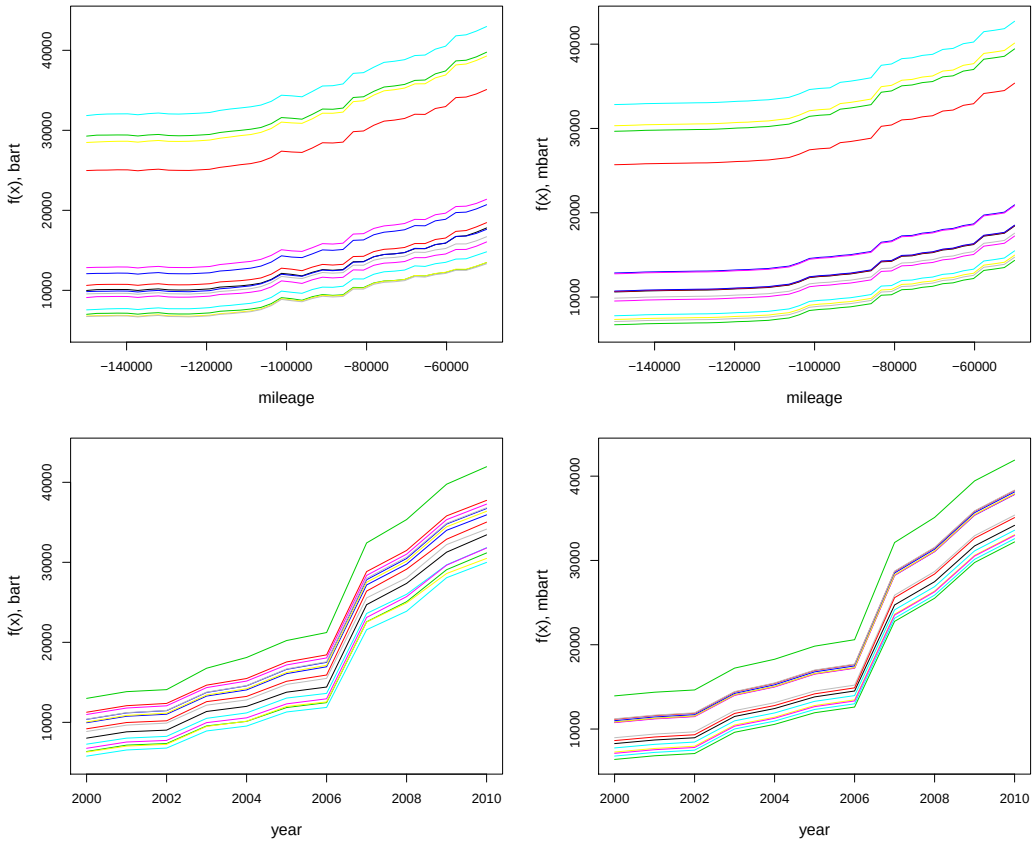

Figure 12: *Car price example. Conditional effects of mileage (top panels) and year (bottom panels) for BART (left panels) and mBART (right panels).*

The conditional effects for `featureCount` and `isOneOwner` are in Figure 13. The conditional effect plot for `featureCount` is quite striking. The monotonic constraint results in a flat lining of the plot dramatically indicating the absence of an effect (contrary to the `***` and t-value of 3.346 in the `R` multiple regression output). After we obtained these results, we checked back with source of the data and found that we had misunderstood the variable `featureCount` and in fact, there is no reason to expect it to be predictive of the car prices! It measures web activity of a shopper and not a feature of the actual car.

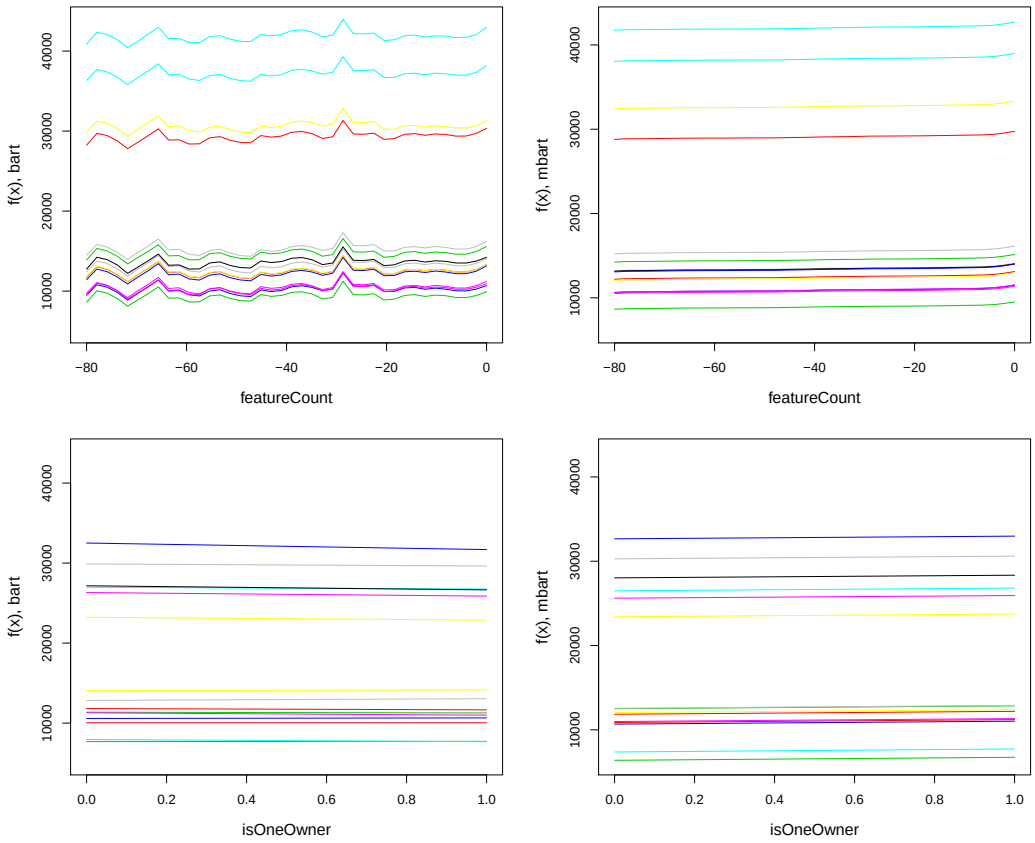

Figure 13: *Car price example. Conditional effects of featureCount (top panels) and isOneOwner (bottom panels) for BART (left panels) and mBART (right panels).*

Figure 14 plots the bivariate fitted surface for expected price as a function of `mileage` and `year` for fixed values of `featureCount` and `isOneOwner`. The BART fit is on the left and the mBART fit is on the right. While similar, the mBART fit is smoother.

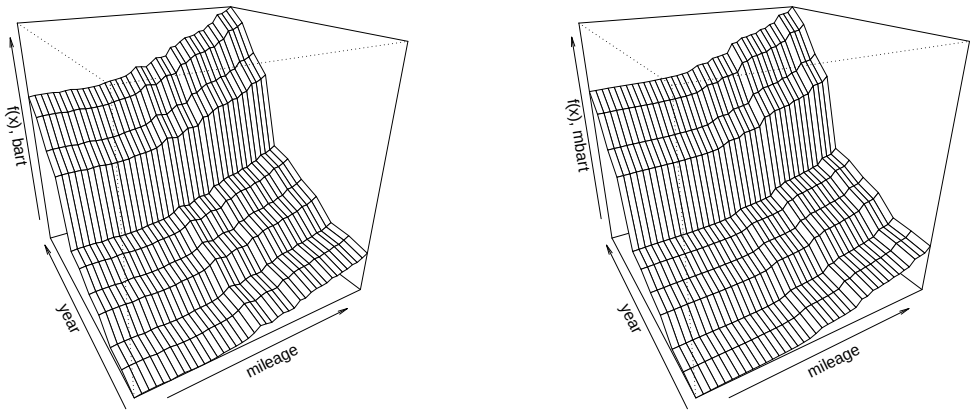

Figure 14: *Car price example. Bivariate plot of fitted price vs. mileage and year. BART (left panel), mBART (right panel).*

We conducted a simple out-of-sample experiment to check for over-fitting: 200 times we randomly selected 75% of the data to be in-sample and predicted the remaining 25% of the y values given their x values using linear regression, BART, and mBART. The out-of-sample root mean square errors are reported in Figure 15. BART and mBART give similar results and both are dramatically better than the linear predictions.

Our results indicate that (i) the monotonicity is indeed reasonable, (ii) we can can get nice monotonic fits, and (iii) only `mileage` and `year` matter.

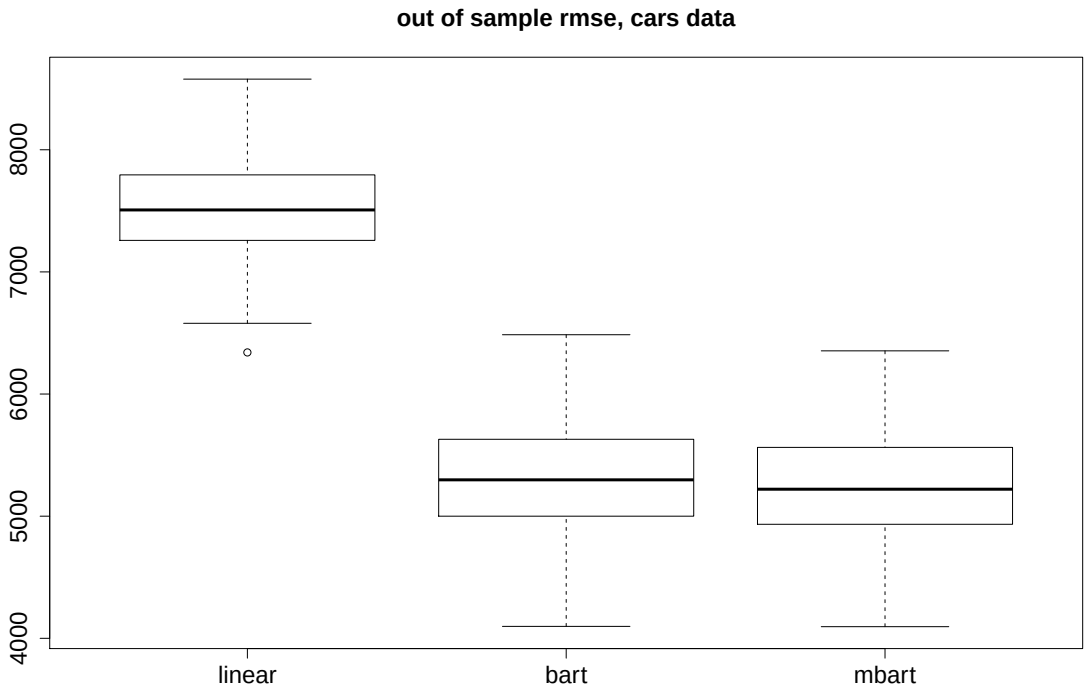

Figure 15: *Car price example. Out-of-sample results.*

5.5 The Stock Returns Example

An important and heavily studied problem is the predictability of stock market returns. Can we measure characteristics of a firm (x) that can help us predict a future return (y)? The data are monthly and the x's are measured the previous month so that the relationship being studied is predictive. Our y is actually *excess* return, the difference between the return for a firm and the average return for firms that month. While still very useful in practice, predicting the excess return is easier than predicting the whole return.

Often in predictability studies predictive models are fit for each month and then rolling windows of months are considered. In this section we pick one month (at random) and consider the fitting of a model for that particular month. We picked December, 1981 from a much larger data set of 594 months. Since the modeling is done for each month, it makes sense to focus on a particular month to see how different approaches might work. However, the predictive model uncovered is exploratory rather than reflective of the actual predictive mechanism used in practice, which is based on averaging rolling fits.

We used four predictive variables in x. `logme`: market equity (logged), `r1`: return, `gpat`: gross profitability ((sales - cost of goods sold) / total assets), and `logag`: growth in total assets (logged). Remember, x is lagged. We have observations on 1,531 firms. Although log transformations of predictors is unnecessary for BART and mBART, these transformations facilitate comparisons with linear regression.

Logged market equity, previous return, and logged growth in total assets are all multiplied by -1 to give monotone increasing relationships. Below are the results of multiple linear regression:

```
            Estimate Std. Error t value Pr(>|t|)
(Intercept) 0.028895   0.010052   2.875  0.00410 **
logme       0.004461   0.001626   2.744  0.00614 **
r1          0.063310   0.020397   3.104  0.00195 **
gpat        0.035634   0.007428   4.797 1.77e-06 ***
logag       0.080160   0.010715   7.481 1.24e-13 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

Residual standard error: 0.07285 on 1526 degrees of freedom
Multiple R-squared:  0.05767,Adjusted R-squared:  0.0552
F-statistic: 23.35 on 4 and 1526 DF,  p-value: < 2.2e-16
```

From the multiple regression output we see that after flipping three of the x variables we do indeed get positive coefficient estimates. It is widely believed that larger firms are less risky

and hence generate lower returns. Hence, the monotonicity for `logme` is strongly motivated. For the other variables, the story is less simple. One might think a high previous return would lead to a high current return (giving a negative sign in the regression since we multiplied by -1). However a tendency for "short term reversals" has been found. Intuitively, gross profitability should be positively related to returns as in our regression. The monotonicity of the `logag` effect is less clear and, indeed, the sign of the regression coefficient can vary from month to month. However, financial theory suggests that if we interpret our x's as representative of underlying factors, the direction of an effect can vary month to month, but we still expect the effect to be monotonic within a given month across a set of firms.

The R^2 in the multiple regression is less than 6%, indicating a very low signal to noise ratio. Bias-variance considerations suggest that only the simplest models can be used to predict since fitting complex models with such a low signal is not feasible. This gives us a strong motivation for examining the fit of mBART. Perhaps using mBART allows us to be more flexible than a simple linear approach without running the risks associated with an unconstrained fit given the low signal.

Figure 16 displays fits from BART, mBART, and a linear regression (using the same layout as in our previous examples). The top left plot shows the sequence of σ draws from the BART fit, while the top right plot shows the sequence of σ draws from mBART. In each plot, a solid horizontal line is drawn at the least squares estimate of σ. The σ draws from the BART fit tend to be smaller than the least squares estimate while the least squares estimate is right at the center of the mBART fits. The monotonicity constraint has pulled the BART fit back so that overall, it is more comparable to the linear fit.

The lower left panel of Figure 16 plots the BART fits versus the mBART fits and the lower left panel plots the linear fits versus the mBART fits. Given the very low signal, it is notable that all three methods pick up similar fits. The mBART fit appears to be more like the linear fit than the BART fit.

Figures 17 and 18 display the conditional effects using the same format as in our previous two examples. The mBART fits are much more appealing. They are quite close to linear (especially for `r1`) but there is some suggestion of nonlinearity for three of the variables.

Figure 19 plots the fitted cross section of expected returns against `r1` and `logme`. The unconstrained BART fit seems quite absurd while the mBART fit suggests some nonlinearity, but also leaves open the possibility that it is close enough to linear for prediction given the high noise level.

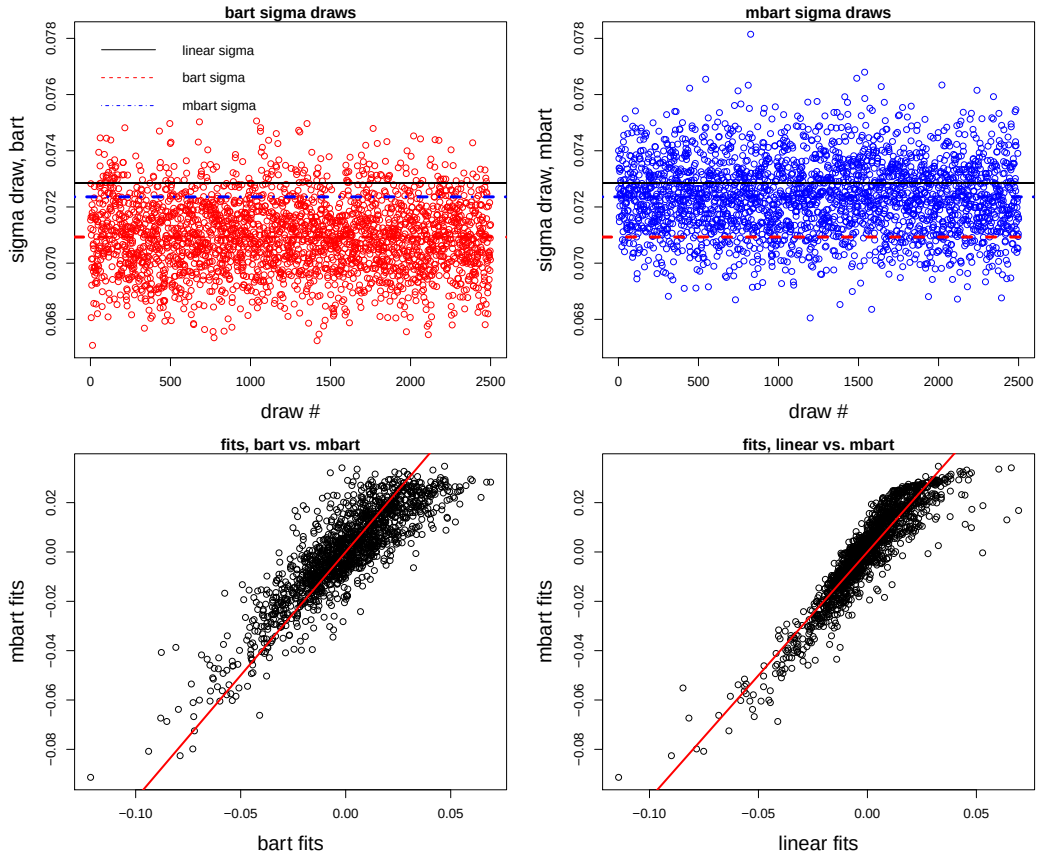

Figure 16: *Returns example. Top row: σ draws from BART (left panel), σ draws from mBART (right panel). Solid horizontal line at least squares estimate of σ. Bottom row: BART versus mBART (left panel) and linear fits versus mBART (right panel).*

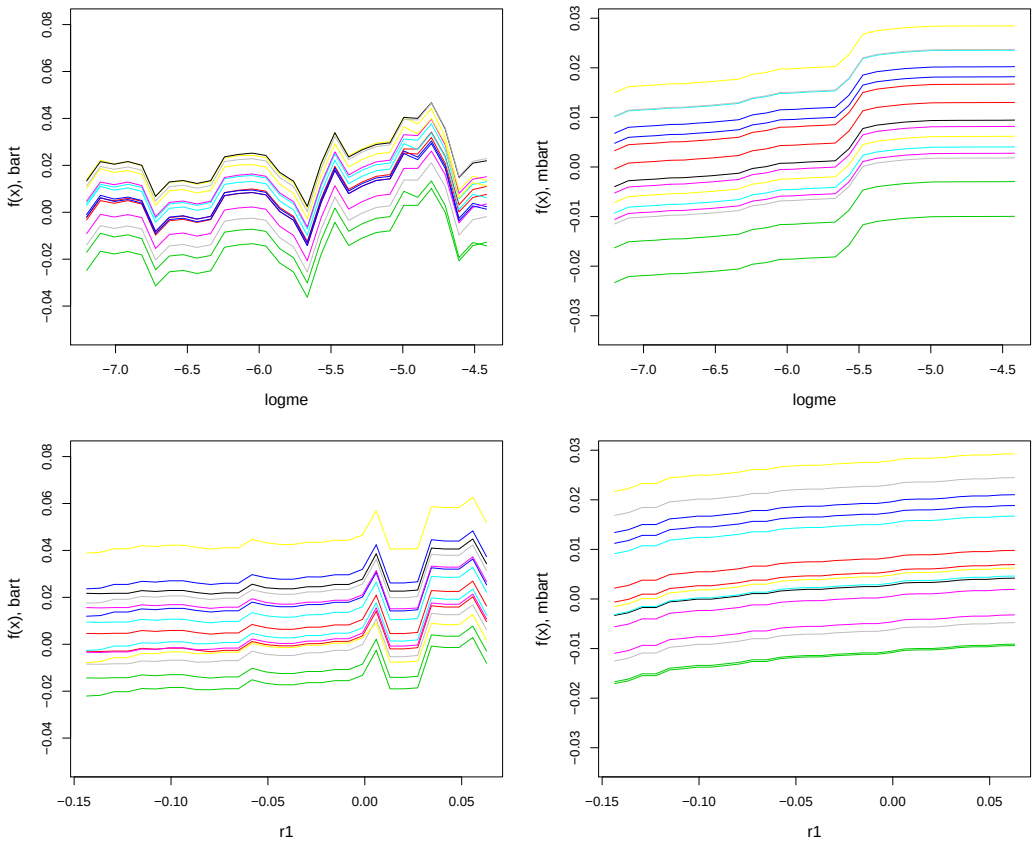

Figure 17: *Returns example. Conditional effects of logme (top panels) and r1 (bottom panels) for BART (left panels) and mBART (right panels).*

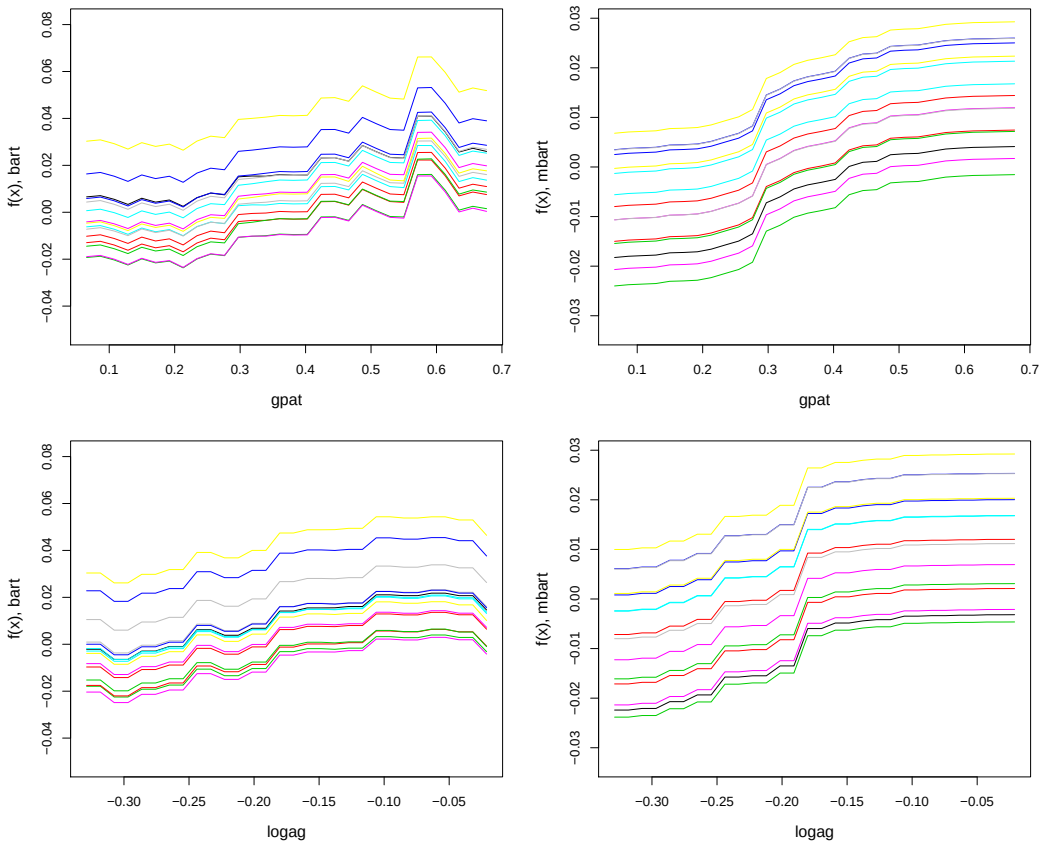

Figure 18: *Returns example. Conditional effects of gpat (top panels) and logag (bottom panels) for BART (left panels) and mBART (right panels).*

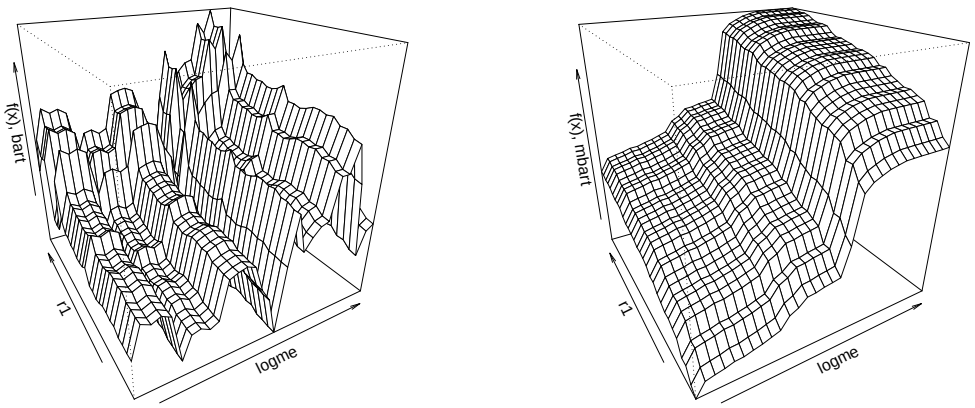

Figure 19: *Returns example. Bivariate plot of fitted expected return vs. logme and r1. BART (left panel), mBART (right panel).*

To get some feeling for out-of-sample predictability, we performed a "stylized" out-of-sample experiment using the same setup as we used for the used cars example in the previous section. We call this "stylized" because it is unrealistic to observe returns from 75% of the firms and then predict the rest. However, this gives a sense for how the procedures work in the particular month. Out-of-sample RMSE's are displayed in Figure 20. We see that mBART and the linear fit give very similar results while BART is slightly worse. Again, given the very low signal to noise ratio, we do not expect to be able to detect large differences.

Our main point is that if you want to consider something more flexible than linear, and interpret the fits on a monthly basis, Figure 19 shows that BART does a poor job, while mBART can give plausible nonlinear results.

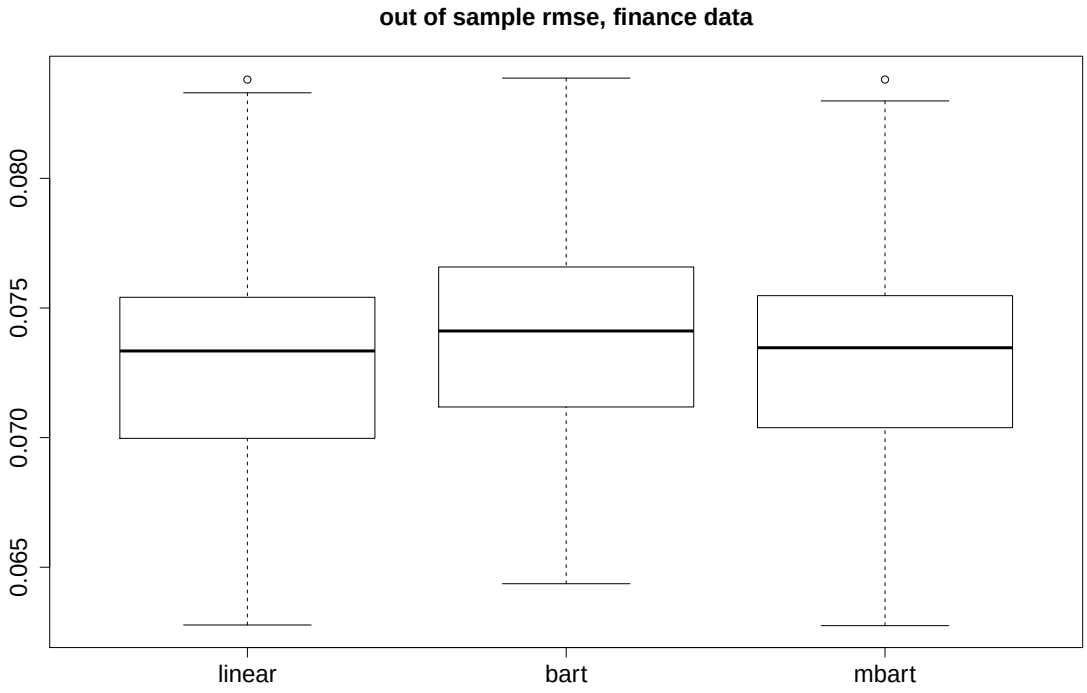

Figure 20: *Returns example. Out-of-sample RMSE. mBART is comparable to linear, while BART is worse.*

6 Conclusion

When prior information is available, we should use it. Information about whether an increase in an explanatory x is likely to increase or decrease the dependent variable Y on average is often available. For example, we do not expect a smaller house to sell for more, all other things being equal. Even in the relatively simple context of linear modeling, prior information about the signs of the coefficients may be helpful.

In the context of highly flexible nonlinear models such as BART, the potential benefits of including prior information regarding monotonicity are even greater. By sensibly constraining our search in the vast space of high-dimensional non-linear functions we ease the burden of search and inference. Our resulting estimates are smoother and more interpretable. We see this in all our examples (Figures 17, 18, and 19 are particularly dramatic). The additional information provides a form of regularization that may lead to better out-of-sample prediction (see Figure 7). Injection of additional prior information can affect both our posterior uncertainty (Figure 5) and the sensitivity of our results to other aspects of our prior specification (Figure 6).

A key advantage of our approach is that it works for high-dimensional x whereas other monotonic approaches become difficult in this case. The simple observation that the sum of monotonic functions is monotonic reduces our problem to that of monotonic inference for a single tree. Since we seek the stochastic search exploration of the posterior uncertainty provided by Markov Chain Monte Carlo algorithms, we adapt the BART MCMC algorithm. This turns out to be non-trivial because a key step in the BART algorithm uses the conditional conjugacy of the prior setup to enable the integration of bottom node mean parameters. This conditional conjugacy is not available in the constrained case. In a fundamental rethinking of the BART algorithm, we make local moves conditional on the rest of the tree including the current values of bottom node mean parameters. This conditioning allows us to locally impose the monotonicity in a simple way and integrate parameters numerically without dramatically slowing the time per iteration.

In some examples, the data are sufficiently informative so that the injection of the additional information regarding montonicity makes little difference and BART results are similar to mBART results. However, when the noise level is relatively high, the additional information can be useful as in Sections 5.2 and 5.5. In all examples, given the obvious appeal of monotonicity, the ability to obtain strictly monotonic estimates provides more interpretable and appealing results.

All data and code used in the paper is available at:
https://bitbucket.org/remcc/mbart as a git repo.
That is, git clone https://bitbucket.org/remcc/mbart will get everthing and then check out README.txt.

References

Azzalini, A. (1985). A class of distributions which includes the normal ones. *Scand. J. Statist.* **12** 171-178.

Barlow, R.E., Bartholomew, D., Bremner, J.M. and Brunk, H.D. (1972). *Statistical Inference Under Order Restrictions: Theory and Application of Isotonic Regression.* New York: Wiley.

Bleich, J., Kapelner, A., George, E.I. and Jensen, S.T. (2014). Variable selection for BART: An application to gene regulation. *Ann. Appl. Stat.* **8** 1750-1781.

Chipman, H., George, E.I. and McCulloch, R.E. (1998). Bayesian CART model search (with discussion and a rejoinder by the authors). *J. Amer. Statist. Assoc.* **93** 935-960.

Chipman, H., George, E.I. and McCulloch, R.E. (2010). BART: Bayesian additive regression trees. *Ann. Appl. Stat.* **4** 266-298.

Chipman, H., George, E.I. and McCulloch, R.E. (2013). Bayesian Regression Structure Discovery. *In Bayesian Theory and Applications*, (Eds, P. Damien, P. Dellaportas, N. Polson, D. Stephens), Oxford University Press, USA.

Holmes, C.C. and Heard, N.A. (2003). Generalized monotonic regression using random change points. *Statist. Med.* **22** 623-638.

Kapelner, A. and Bleich, J. (2016). bartMachine: Machine learning with Bayesian additive regression trees. *J. Stat. Softw.* **70** 1-40.

Kong, M. and Eubank, R.L. (2006). Monotone smoothing with application to dose-response curve. *Commun. Statist. B-Simul.* **35** 991-1004.

Lavine, M. and Mockus, A. (1995). A nonparametric Bayes method for isotonic regression. *J. Stat. Plan. Inference*, **46**, 235-248.

Lin, L. and Dunson, D.B. (2014). Bayesian monotone regression using Gaussian process projection. *Biometrika* **101** 303-317.

Mammen, E. (1991). Estimating a smooth monotone regression function. *Ann. Statist.* **19** 724-740.

Neelon, B. and Dunson, D.B. (2004). Bayesian isotonic regression and trend analysis. *Biometrics* **60** 177-191.

Ramsay, J.O. (1998). Estimating smooth monotone functions. *J. R. Statist. Soc. B* **60** 365-375.

Saarela, O. and Arjas, E. (2011). A method for Bayesian monotonic multiple regression. *Scand. J. Statist.* **38** 499-513.

Shively, T.S., Sager T.W. and Walker, S.G. (2009). A Bayesian approach to nonparametric monotone function estimation. *J. R. Statist. Soc. B* **71** 159-175.

Shively, T.S., Walker, S.G. and Damien, P. (2011). Nonparametric function estimation subject to monotonicity, convexity and other shape constraints. *J. Econometrics* **161**, 166-81.